

PHYSICAL ACOUSTICS SUMMER SCHOOL



19970421 001

ASTILOMAR CONFERENCE CENTER

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

DTIC QUALITY INSURED 3

VOLUME I: TRANSCRIPTS

1996 PHYSICAL ACOUSTICS SUMMER SCHOOL

VOLUME I: TRANSCRIPTS

This work relates to Department of Navy Grant N00014-96-1-0033 issued by the Office of Naval Research. The United States Government has a royalty-free license throughout the world in all copyrightable material contained herein.

DTIC QUALITY INSPECTED 3

Copies of this three-volume proceedings can be obtained by contacting: Libby Cauthen, NCPA, University of Mississippi, University, MS 38677; voice: 601-232-5808; fax: 601-232-7494; e-mail: eacauthe@olemiss.edu.

TABLE OF CONTENTS

FORWARD, LOGAN E. HARGROVE.....	1
LIST OF PARTICIPANTS	2
GENERAL BACKGROUND, STEVEN L. GARRETT	5
MOLECULAR ACOUSTICS, HENRY E. BASS	55
RUS AND MATERIALS PHYSICS, ALBERT MIGLIORI.....	88
PERIODIC, RANDOM AND QUASIPERIODIC MEDIA, JULIAN D. MAYNARD	132
SONOLUMINESCENCE, ANTHONY A. ATCHLEY.....	177
FUNDAMENTALS AND APPLICATIONS OF NONLINEAR ACOUSTICS, MARK F. HAMILTON	210
ATMOSPHERIC ACOUSTICS, RICHARD RASPET	254
SENSOR PHYSICS: SIGNALS AND NOISE, THOMAS B. GABRIELSON	286
THERMOACOUSTIC ENGINES AND REFRIGERATORS, GREGORY W. SWIFT	328
REPORT DOCUMENTATION PAGE.....	368

1996 PHYSICAL ACOUSTICS SUMMER SCHOOL

FORWARD

These are the Proceedings of the 1996 Physical Acoustics Summer School (PASS 96). The lectures were recorded and the verbatim transcripts were subsequently edited by the authors for publication here.

The Office of Naval Research (ONR), in cooperation with the Acoustical Society of America (ASA) and the National Center for Physical Acoustics (NCPA) sponsored PASS 96, June 21-28, 1996 at the Asilomar Conference Center in Pacific Grove, California, the site of the first and subsequent Summer Schools, PASS 92 and PASS 94. Participation in the three Summer Schools was limited to a total of 50 which included students, lecturers, and discussion leaders.

The purpose of these Summer Schools is to bring graduate students, distinguished lecturers, and discussion leaders together to discuss a wide variety of subjects in physical acoustics. This gives the students the opportunity to meet experts and talk about topics most students ordinarily wouldn't encounter at their own colleges and universities. The focus was on graduate students and academic participants. Approximately half of the participants have been advanced graduate students in physical acoustics.

The Summer Schools have their beginning in an ONR Principal Investigators meeting in 1988 where it was decided that the best investment of this kind would be in a Summer School focused on graduate students. In 1990, an informal Summer School for students of ONR Principal Investigators and some invited guests was held as part of a Principal Investigator's meeting at the Naval Postgraduate School, Monterey, California, and this controlled experiment set the pattern and influenced the site selection for the subsequent PASS 92, PASS 94 and PASS 96. We hope that the Summer Schools will continue as biennial events with the high standards and wonderful success we have thus far enjoyed. This has happened and will happen again because everyone involved does their best, and for this I say, "Thank You!!!"

LOGAN E. HARGROVE
OFFICE OF NAVAL RESEARCH

1996 PHYSICAL ACOUSTICS SUMMER SCHOOL

LIST OF PARTICIPANTS

———LECTURERS———

DR. ANTHONY A. ATCHLEY
NAVAL POSTGRADUATE SCHOOL
DEPARTMENT OF PHYSICS
CODE PH/AY
MONTEREY, CALIFORNIA 93943

DR. HENRY E. BASS
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MISSISSIPPI 38677

DR. THOMAS B. GABRIELSON
APPLIED RESEARCH LABORATORY
PENNSYLVANIA STATE UNIVERSITY
P.O. BOX 30
STATE COLLEGE, PA 16804-0030

DR. STEVEN L. GARRETT
PENNSYLVANIA STATE UNIVERSITY
GRADUATE PROGRAM IN ACOUSTICS
P.O. BOX 30
STATE COLLEGE, PENNSYLVANIA 16804

DR. MARK. F. HAMILTON
UNIVERSITY OF TEXAS
DEPT. OF MECHANICAL ENGINEERING
AUSTIN, TEXAS 78712-1063

DR. ROBERT M. KEOLIAN
NAVAL POSTGRADUATE SCHOOL
PHYSICS DEPARTMENT
CODE PH/KN
MONTEREY, CALIFORNIA 93943

DR. JULIAN D. MAYNARD, JR.
PENNSYLVANIA STATE UNIVERSITY
DEPARTMENT OF PHYSICS
104 DAVEY LABORATORY
UNIVERSITY PARK, PENNSYLVANIA
16802

DR. ALBERT MIGLIORI
LOS ALAMOS NATIONAL LABORATORY
P.O. BOX 1663
MSK764
LOS ALAMOS, NEW MEXICO 87501

DR. RICHARD RASPET
UNIVERSITY OF MISSISSIPPI
DEPARTMENT OF PHYSICS & ASTRONOMY
UNIVERSITY, MISSISSIPPI 38677

DR. GREGORY W. SWIFT
LOS ALAMOS NATIONAL LABORATORY
MAIL STOP K764
LOS ALAMOS, NEW MEXICO 87501

——DISCUSSION LEADERS——

DR. DAVID A. BROWN
UNIV. OF MASS. @ DARTMOUTH
DEPT. OF ELECTRICAL AND COMPUTER
ENGINEERING
285 OLD WESTPORT ROAD
NORTH DARTMOUTH, MASSACHUSETTS
02747-2300

DR. EDWARD F. CAROME
JOHN CARROLL UNIVERSITY
PHYSICS DEPARTMENT
CLEVELAND, OHIO 44118

DR. JAMES P. CHAMBERS
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MISSISSIPPI 38677

DR. DAVID L. GARDNER
LOS ALAMOS NATIONAL LABORATORY
MS K764
LOS ALAMOS, NEW MEXICO 87545

DR. LOGAN E. HARGROVE
OFFICE OF NAVAL RESEARCH
ONR 331
800 NORTH QUINCY STREET
ARLINGTON, VIRGINIA 22217-5660

DR. THOMAS J. HOFER
NAVAL POSTGRADUATE SCHOOL
CODE PH/HF
MONTEREY, CALIFORNIA 93943

DR. MOISES LEVY
UNIVERSITY OF WISCONSIN-MILWAUKEE
PHYSICS DEPARTMENT
MILWAUKEE, WISCONSIN 53201

DR. PHILIP L. MARSTON
WASHINGTON STATE UNIVERSITY
DEPARTMENT OF PHYSICS
PULLMAN, WASHINGTON 99164-2814

DR. RICHARD STERN
PENNSYLVANIA STATE UNIVERSITY
APPLIED RESEARCH LABORATORY
P.O. BOX 30
STATE COLLEGE, PENNSYLVANIA 16804

DR. DAVID B. THIESSEN
WASHINGTON STATE UNIVERSITY
PHYSICS DEPARTMENT
PULLMAN, WASHINGTON 99164-2814

————STUDENTS————

MR. BRIAN O. BENNETT
PENNSYLVANIA STATE UNIVERSITY
P.O. BOX 924
STATE COLLEGE, PENNSYLVANIA 16804

MR. REH-LIN CHEN
PENNSYLVANIA STATE UNIVERSITY
445 WAUPELANI DRIVE
APARTMENT F26
STATE COLLEGE, PENNSYLVANIA 16801

MR. XIAOHUI CHEN
YALE UNIVERSITY
9 HILLHOUSE AVENUE
NEW HAVEN, CONNECTICUT 06520

MR. BARRY DOUST
APPLIED RESEARCH LABS
PENNSYLVANIA STATE UNIVERSITY
P.O. BOX 30
STATE COLLEGE, PENNSYLVANIA

MS. ANAT GRANT
PENNSYLVANIA STATE UNIVERSITY
GRADUATE PROGRAM IN ACOUSTICS
P.O. BOX 30
STATE COLLEGE, PENNSYLVANIA 16804

QUISHUANG GUO
UNIV. OF MASS. AT DARTMOUTH
22 SHARON AVENUE
NORTH DARTMOUTH, MASSACHUSETTS
02747

MR. IBRAHIM HALLAJ
UNIVERSITY OF WASHINGTON
APPLIED PHYSICS LABORATORY
1013 NE 40TH STREET
SEATTLE, WASHINGTON 98105-6698

MR. BRIAN T. HEFNER
WASHINGTON STATE UNIVERSITY
1045 HARVEY ROAD, #A
PULLMAN, WASHINGTON 99163

MR. MARK C. KELLY
UNIVERSITY OF MISSISSIPPI
NCPA
UNIVERSITY, MS 38677

MR. JEFFREY A. KETTERLING
YALE UNIVERSITY
279 HUMPHREY STREET
NEW HAVEN, CONNECTICUT 06511

BRIAN J. LANDSBERGER
THE UNIVERSITY OF TEXAS @ AUSTIN
1900 MISTYWOOD DRIVE
AUSTIN, TEXAS 78746

MR. ANTON I. LAVRENTYEV
OHIO STATE UNIVERSITY
190 WEST 19TH AVENUE
COLUMBUS, OHIO 43210

MR. WENSEN LIU
SAN DIEGO STATE UNIVERSITY
COLLEGE OF ENGINEERING
5500 CAMPANILE DRIVE
SAN DIEGO, CALIFORNIA 92182

XIAORONG LU
CORNELL UNIVERSITY
DEPT. OF THEORETICAL AND APPLIED
MECHANICS
212 KIMBALL HALL
ITHACA, NEW YORK 14853

MR. MARK J. MARR-LYON
WASHINGTON STATE UNIVERSITY
100 NW TERRE VIEW #A3
PULLMAN, WASHINGTON 99163

MR. PATRICE MASSON
SHERBROOKE UNIVERSITY
DEPARTMENT OF MECHANICAL
ENGINEERING
SHERBROOKE, P.Q., J1K 2R1
CANADA

MR. SCOT F. MORSE
WASHINGTON STATE UNIVERSITY
1015 NE DUNCAN LANE
PULLMAN, WASHINGTON 99163

MR. MATTHEW E. POESE
PENNSYLVANIA STATE UNIVERSITY
GRADUATE PROGRAM IN ACOUSTICS
P.O. BOX 30
STATE COLLEGE, PENNSYLVANIA 16804

MR. WAYNE E. PRATHER
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MISSISSIPPI 38677

JIN-HYUN SO
OHIO UNIVERSITY
DEPT. OF PHYSICS AND ASTRONOMY
ATHENS, OHIO 45701

MR. DORU VELEA
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MISSISSIPPI 38677

MR. YUNLONG WANG
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MISSISSIPPI 38677

MR. MARTIN J. WETZEL
JOHNS HOPKINS UNIVERSITY
2913 GUILFORD AVENUE
BALTIMORE, MARYLAND 21218

MR. PHILLIP WHITE
PENNSYLVANIA STATE UNIVERSITY
619 S. PUGH STREET., #307
STATE COLLEGE, PENNSYLVANIA 16801

MR. ANIRUDDHA S. WORLIKAR
JOHNS HOPKINS UNIVERSITY
122 LATROBE HALL
3400 N. CHARLES STREET
BALTIMORE, MARYLAND 21218

-----STAFF-----

MS. LIBBY CAUTHEN
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MISSISSIPPI 38677

MS. ROMELLE S. MILLION
MILLION REPORTING
P.O. BOX 2419
ALEXANDRIA, VIRGINIA 22301

GENERAL BACKGROUND

STEVEN L. GARRETT
PENNSYLVANIA STATE UNIVERSITY

DR. GARRETT: I was very honored to be asked to give the opening lecture for summer school this year, until I spoke to Anthony's wife, Mary Atchley, who said, "Ah, well, you're the opening act, you're the guy they want to get off the stage quickly so that the real headliners can give their talks." So I will bore you for a few hours while the headliners are getting ready.

(Transparency #1)

The summer schools -- this is, I think, the fourth that I have been to here and, also, I will mention the one in Varenna -- from all we can tell have been very successful. One way to tell that is just being here; you can feel when things are going well and you can tell when science and ideas are getting bounced around.

Another is that there is a questionnaire at the end and people are asked for their suggestions and their comments. Given that historical success, the organizing committee -- Drs. Hargrove, Bass, and Atchley -- are rather conservative. They know a good thing, they have a good thing, and they are not anxious to make dramatic changes, but this time they decided to try something different and that is to have a background lecture that would act to bring the various diversities of educational experiences that people bring with them here and try to not necessarily fill in gaps but provide a common context into which the specialty lectures you will hear over the next week will fit.

It is clearly an experiment. This has never been done before. I do not know whether it is going to be successful or not, but we are going to do it. Instead of having a real lecture at the beginning, we are going to have this general background lecture.

I had to start somewhere, so I figured I would start here: There is no sound in a vacuum. Let us talk about materials. Immediately Professor Bass comes up to me and says that is wrong. "Your first viewgraph is wrong." Tell them why you said it is wrong.

DR. BASS: I do not understand the statement that there is no sound in a vacuum. We all know that in intergalactic space -- [momentary technical difficulties with the recording] -- ways of propagating which are quite observable and very beautiful.

DR. GARRETT: I was just trying to get a start.

(Laughter)

You guys did not hand me a lot here. We will talk about nucleation and bubble formation; I needed a nucleus here.

DR. BASS: If you have a theoretical vacuum, there are vacuum fluctuations where you perturb one wall in a container and other walls in the container will sense that transmission through a transmission even in the theoretically absolute vacuum, so I am not sure that is a correct statement.

DR. GARRETT: Okay, well, let us start with the outline, then.

(Laughter)

(Transparency #2)

This is an experiment. I will hold on this bullet. This is definitely an experiment. But I will yield on the second point and move on to the third point, unless there is some comment about the ability to extract information from vacuum fluctuations in violation of well-known laws of thermodynamics.

When I looked at this stuff, I knew what some of the lecturers were going to be covering and I had a vague idea of what some of the other lecturers were going to be covering, and I had no handle whatsoever on Jay Maynard. I tried to set things up in a way that would give you as much context as possible for the largest number of lecturers.

Two years ago Professor Bass gave the lecture on molecular acoustics and I wanted to add a little bit of balance to that view, so I am going to talk about the two complementary pictures that science has of nature:

1. the phenomenological picture, which asks what are the macroscopic observables and what are the conservation laws that dictate the evolution of those variables;
2. and the microscopic picture, the quantum mechanical picture. What is that fluid actually made of? How do the properties of the individual atoms or molecules and their quantum mechanical states, affect those variables?

I want, in the first part of the lecture, to address this idea of the two complementary models, the phenomenological model and the microscopic, or quantum mechanical, model, that we use to explain acoustic phenomena.

After failing to do that, I will move on to ideal gas thermodynamics. It turns out that I think that will be useful in getting a common language for adiabatic and isothermal processes and for some simple ideas and simple applications of this picture of phenomenology and microscopics.

I have included an application which was in the handouts that were sent to the students but not to the lecturers on acoustic gas analysis and, actually, acoustic pyrometry (measuring temperature with gas), and a few comments about the measurement of fundamental constants using ideal gas thermodynamics.

At that point I will go on from equilibrium processes to nonequilibrium processes, or irreversible processes and transport properties. This is something that is fundamental to our

understanding of acoustics. I think it will be beneficial in several of the lectures, ranging from bubble dynamics to attenuation and relaxation processes.

So I will work on that for a while and then do another application, which is the calculation of resonator quality factors due to these thermoviscous interactions that take place in the boundary between moving gases and motionless solids.

Finally, I would like to talk about isotropic elasticity. Al Migliori is going to be talking about solids. He is going to be talking about resonant ultrasound spectroscopy. I thought I would just make a few comments about the relationship between the modes of simple vibrating objects and the elastic properties of the materials that are doing the vibration.

That is the outline of what I hope to accomplish; we will see how it goes -- as long we remember that I am an experimentalist and this is an experiment.

(Transparency #3)

Before I launch into the material I would like to make some comments that I think Logan made last night, but I would like to repeat them. This is a picture of the summer school that was held in Varenna, Italy, in 1974, at Villa Monastero on Lago di Como (Lake Como), which is right at the Swiss-Italian border.

The Italian Physical Society was given a villa there, Villa Monastero, which was, in my picture, at that age, just what you think an Italian villa should be: fancy, marble statues of guys holding grapes, verandas, and leather-tooled wallpaper -- just really decadent Italian beachfront property.

For six weeks a year the Italian Physical Society had it and they would run the Enrico Fermi Summer School, which was three two-week summer-school classes. About every 10 years this gentleman, Daniel Sette, a very well-known ultrasonic scientist, would come up to the Italian Physical Society to hold a summer school. So in 1964, 1974, and 1984 there were summer schools that were related to acoustics (particularly physical acoustics).

That venue has evaporated as Professor Sette has gone into retirement, but I was at the summer school at Varenna in 1974 and it made a very great impression on me. This is why I am bringing it up with you. Physical acoustics is not a very big field and it is not a very richly endowed field. There is not a lot of money in physical acoustics and there are not a lot of universities that have a critical mass where there is a lot of faculty and a lot of students all concerned with the same problem.

But it is a fascinating field and its practitioners are truly wonderful people -- I will point some of them out in this picture. They would get together to run these summer schools, because there were not big solid-state physics departments in which there were 10 solid-state physicists and \$10 million in research grants, for their graduate students. It was a wonderful thing when I was this guy back here who looks like Charles Manson.

(Laughter)

To see great physicists -- look at the front row. This guy is Joe Hunter. This is Bruce Lindsay. Next to him is Warren Mason. Bruce Lindsay was the godfather of acoustics at Brown University when Brown had an enormous acoustics operation. This is Warren P. Mason, who was basically the chief acoustician at Bell Labs during the heyday at Bell Labs. Daniel Sette, who was a professor in Italy.

This is my esteemed adviser, whom I do not mention by name, because my students always have to stand when his name is mentioned. I just refer to him as "the mentor." I believe this is Klaus Dranesfield -- I am not sure. This is Richard K. Cook. He was from the Bureau of Standards.

In the 1994 proceedings, Professor Bass declared him dead. He is not dead. I spoke to him about a month ago.

(Laughter)

But if you look at the proceedings, Professor Bass was talking about the measurement of low-frequency sound using these capacitive microphones with a F.M. carrier system and he said that the business in those 'phones had dropped off "since he died." The guy is still alive, so those of you who have the proceedings from 1994 should make that correction. He was there and he gave an excellent lecture on infrasound, on listening to the surf breaking in California while you were in Maryland.

This guy here is on the old guys' bench --

DR. BASS: He is not dead, either.

(Laughter)

DR. GARRETT: Right, nobody ever said you were dead -- it is Ed Carome, who is in the back with us. There are a few other people here you might know, you Penn Staters, Ralph Goodman (with hair), John McClerk, a very famous man in surface acoustic wave devices, Taylor Wang, acoustic levitation, who flew in a space shuttle, STS-24.

There are some people there I do not know, some are just graduate students -- Joe Heiserman, who is now a physician in Arizona. He was with me at UCLA in the mentor's laboratory. There was a great group, about the same size as this, in Italy, and we would hear these lectures and I, as a graduate student, was very impressed. I was very impressed by sitting around with these guys and listening to them discuss scientific topics of the day.

In physical acoustics, that, I think, is the tradition, and that is what we are trying to do here. We are trying to give you exposure to a very wide variety of topics that you will not get in your own university. You may get one or two, but you will not see that greater picture of how people interpret acoustics, particularly people who have been doing it for the past 20 or 30 years of their lives.

I hope to be one of these old guys on the bench 22 years from now. Those guys did not have to lecture. We just "osmosed" when they were around. I am hoping that some of you get the message here and you remember that you are the beneficiaries of this very wonderful tradition. You will be called upon in the future to continue that tradition.

Speaking for the consortium of bald scientists, I wanted to put that in historical context and to let you know that we have not collected the real tuition from you students yet. We will collect the real tuition from you guys 20 years from now when you get to do this.

(Transparency #4)

Let us start with the material. There are two basic ways to look at science, and they are complementary ways. It is not that one way is right and the other way is wrong, or that one way is more elegant and the other way is messy. They are both important and they feed each other. The two approaches have to be understood in that way.

Phenomenology, I think, is best summarized by a friend of mine, Andrés Larraza, who says "an acoustician is merely a timid hydrodynamicist." With all due respect to the resident ultrasound solid guys, also we could say merely a timid elasticity specialist.

But remember, acoustics is the science of the subtle. We are talking about very small deviations about equilibrium. It is a very delicate and subtle science, with the exception of what Drs. Hamilton and Swift are going to talk about.

On the other hand, there is more than just simply a hydrodynamic description of the matter in which sound propagates. There are microscopic models. Richard Feynman said, in the first chapter of his book, *Lectures in Physics*, "If, in some cataclysm, all of the scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words?" Professor Feynman answers his own question: "I believe it is the atomic hypothesis; that all things are made of atoms."

So our understanding, which is relatively new in the history of science -- remember that Boltzmann committed suicide over this at the turn of the century -- this picture, that all matter is made of atoms and the motion of these atoms and the interaction of these atoms, and the behavior of atoms is responsible for the way macroscopic material behaves is very, very important.

For those of you who do not look at things this way, let me define what I mean by a phenomenological model.

(Transparency #5)

In a phenomenological description, you do not worry about what stuff is made of and what the interactions of the constituents are; you treat the universe as being an undifferentiated continuum, a smooth liquid of some kind. The first phenomenological question which must be

answered to make progress is how many variables are required to provide a complete description of the system?

To give you an example, a static, homogeneous, isotropic, single component fluid, a single-phase liquid, requires two variables. It requires a mechanical variable like pressure, density, or volume, and a thermal variable like temperature or entropy. The two variables completely describe the behavior of a homogeneous, isotropic, single-component fluid, if it is at rest.

If it is in motion, you require five variables. You require the two thermodynamic variables for the fluid at rest plus the three components of velocity. So for a single-component fluid, five variables form a complete description.

This, of course, is a major leap. We know that in this much stuff there are 10^{23} atoms and all this action going on. The fact is that we can successfully describe, under certain constraints, the behavior of that system with only two variables and, if it is moving, only five variables.

More complex systems require more variables. This should not come as a major surprise. If we look at something like a superfluid, that is, a fluid whose behavior is determined by the laws of quantum, rather than classical mechanics, then eight variables are required. You need two thermodynamic variables and two vector velocity fields. The quantum fluid can be described as having a normal component, with viscosity, and a superfluid component, without viscosity and, therefore, eight variables for a complete description.

There is a microscopic picture that involves phonons and involves rotons. It requires knowledge of the ground state wave-function and Bose-Einstein condensation and many other details. But in fact, you can blow all those off and eight variables will completely describe all the phenomenon that have ever been observed in superfluid flows.

STUDENT: Is that true for superfluid helium-3?

DR. GARRETT: Right. Superfluid helium-3 has other degrees-of-freedom and hence more macroscopic variables are required. It has a magnetic moment, so it is sort of a quantum plasma in some sense. Plasma physics has charged particles, so there is species equilibrium and concentration of ions and the number gets larger.

So the fundamental question is how many variables form a complete description of a given system? Once we have that, we can close the description in the mathematical sense. That is, if we have five variables and we have five equations for those five variables, then the system has a solution. That is just an algebraic truth. We call such a system "closed".

If we have a number of equations that is equal to the number of variables, either by applying conservation laws or something that we know about the equation-of-state, then the description is closed. That does not make it right. There is only one thing that makes a description right in science: agreement with experiment. Got that?

It does not matter whether you have a fancy academic chair. It does not matter whether you are the richest or the poorest guy on the planet. It does not matter how many pieces of H-P gear you have in your lab. What matters is does it agree with experiment? If it does not agree with experiment, then you have to try a different number of variables and find some other conservation laws.

You have a number of variables. You have a number of equations. That forms a closed description. Whether that description is or is not a good description of nature is answered only by comparison to experiment.

This phenomenological approach has been recognized as being extraordinarily important. The quote I take here is from Albert Einstein. "Thermodynamics" (which is a phenomenological model) "is the true testing ground of physical theory since it is model independent." The results do not depend on whether there are hard little spheres in your gas. They do not depend on whether or not those spheres have internal degrees-of-freedom. They do not depend on whether or not there are gremlins riding sky rockets or angels dancing on little pins. It does not matter what was the internal model you use to guide your intuition; the thermodynamic results are an inevitable consequence of the phenomenology. Therefore, it is the true testing ground of physical theory, according to Einstein. If anyone wants to challenge that, you know who to talk to.

(Transparency #6)

I would like to get started with the simplest possible phenomenological theory, the ideal gas law: $pV = nRT$ (eq. 1). If I insult you with the simplicity of my approach, I am sorry, but I wanted to start from the simplest place I could start, rather than introduce exciting new material. This you saw in high school chemistry, right, $pV = nRT$, it is the ideal gas law; $pV = \text{a constant}$ is Boyle's law.

At constant, V , p is proportional to T , that is known as Charles' law. If the people in electricity and magnetism had the same kind of approach, then Ohm's law would have three names. It is one law, but Charles made ideal gas thermometers out of this thing and Boyle did something else, and they each got to put their name on it, but there is one law.

n is the number of moles, M is the atomic or molecular weight of the species -- for mixtures that may be a weighted average -- R is the universal gas constant (I will have more to say about that). We can write this in an extensive form, as I have done here; that is, if I double the number of moles, I double the volume.

Or we can write it in an intensive form (eq. 2). That is, I can write n , the number of moles, as the mass of gas divided by the molecular weight. I can bring V over to this side and I get p is equal to $\rho \text{ times } RT/M$, where ρ does not depend on how big the system is and p does not depend on how big the system is.

That is a phenomenological model, too, but if I specify p and T , then I know what ρ is. This happens to be a mistake of not knowing my word processor well enough. This is supposed to be a ρ and it is something else. We will get that corrected when the proceedings come out.

DR. SWIFT: You do not have a correct value for the gas constant. I mean, everybody should know the Moldover route to the gas constant. You have got the last digit wrong.

DR. GARRETT: Thank you, Dr. Swift. He is absolutely correct. This will not be the last time. I took that number off the NBS before the Moldover measurement. Thank you. We will get that corrected. What is the right digit?

DR. SWIFT: Two.

DR. GARRETT: Two, not a 3?

DR. SWIFT: Yes.

DR. GARRETT: Sorry. It is an important point that Dr. Swift brings out.

(Laughter)

I know at least one person here knows what the worst known constant in physics is.

DR. MARSTON: G.

DR. GARRETT: Right, big G, Newton's big G. The second worst known constant is Boltzmann's constant. It was redetermined, as Dr. Swift pointed out, in an acoustics experiment that I will have something to say about later on. In fact, our understanding of acoustics is so good that it forms the best determination of the Boltzmann's constant, or the gas constant. So it was important that I missed that fourth digit.

Lord Rayleigh got a Nobel Prize for the fourth digit. Is everybody aware of that? Lord Rayleigh was a very careful experimentalist. He determined the density of nitrogen prepared in two different ways. He determined the density of nitrogen prepared by extracting gases from the atmosphere, extracting the oxygen and the other gases from the atmosphere, and got one value that differed by a part in 1,000 from the value that he got when he chemically prepared nitrogen. Because he was such a good experimentalist, he knew that was a larger discrepancy than his experimental error and went on to discover that there was argon.

Argon was a constituent that is not extracted when you do a chemical reduction of the atmosphere. If you react all the stuff out chemically, you do not get the argon out, because argon is inert. Twenty years earlier, a guy named John Dewar, who invented the vacuum vessel that you use for storing cryogenic liquids and probably was related to the Scotch manufacturer --

(Laughter)

Anyway, he would liquefy air and there would be this white scum on the top which he would skim off: solid argon. If he would have thought a little bit about why there was scum on top of his liquid air, he would have had the Noble Prize 20 years earlier. So do not scoff when

Professor Swift chastises me for an error in the fourth digit. If you know what you are doing, that digit can be very important. Sorry for the digression.

This is a microscopic model of a gas. These little things are hard spheres and they bounce around and their interactions are governed by Newton's law. We did not make any of those assumptions when we did phenomenology. We said how many variables and what are the conservation terms.

But now we are going to say there is something going on in this box. What is in this box are little hard spheres. The little hard spheres are much smaller than the typical space between those hard spheres and they are moving around. Each one has a velocity. I have chosen one here; it has a velocity in this direction. It has a velocity in the x-direction, a component in the x-direction, in the z-direction, and in the y-direction, the vertical direction in this picture.

We will postulate that these little hard spheres bouncing around in this box suffer only elastic collisions. When this guy hits this wall, which is perpendicular to the x direction, called the x wall, it changes its momentum because it reverses its direction and does not change its speed. Therefore, there is a change in momentum that is $2mv_x$. It had mv_x coming into the wall, it had mv_x going out of the wall. This is high school stuff.

There is also a y- and a z-component, but that is discussed on slide 18. When we talk about viscosity, we will come back to that.

What is the collision rate? We know it collides with this wall and goes over to this wall and comes back and comes back and collides with that wall, again. Since we have broken this into Cartesian coordinates, we know that the time in-between collisions of this molecule with this wall is twice the length divided by v_x .

By Newton's second law, we know that the change in momentum per unit time, the rate-of-change of momentum, is the force in the x-direction, that is, on the wall it was normal, it was in the x-direction. We can then plug in and say that the force on this wall per particle is mv_x^2/L .

The pressure on the wall is the force on the wall divided by the area of the wall. That force, divided by $L_y L_z$, the length in the y-direction times the length in the z-direction, that is the area, is, therefore, mv_x^2 over the volume of the box.

It is a microscopic model. We are talking specifically about the dynamical interactions of the particles among themselves and their interactions with the walls of the box.

(Transparency #7)

At that point we add another assumption about the world and about the way particles operate. Here we introduce a very important concept for Tom Gabrielson's lecture: the equipartition theorem. The equipartition theorem is actually proven quite nicely in Feynman's book, but we are just going to take it to be the last vestige of communism. The Berlin Wall is

gone, the Soviet Union has disbanded more or less, but when it comes to molecules, each gets its fair share: "from each according to its energy, to each according to its energy."

If you have got a molecule that is moving around very rapidly in any given direction, it collides into other slower ones until that extra motion is equally shared. When it is all equally shared, we call that equilibrium. We will stick with just the quadratic degrees-of-freedom; that is, their energy is quadratic combinations of their variables. For example, kinetic energy is proportional to velocity squared. Potential energy is proportional to displacement squared. They each get their fair share of the thermal energy and that share is one-half of $k_{\text{Boltzmann}}$ times the absolute temperature.

The temperature, therefore, is related to the kinetic energy. The velocity is the Pythagorean sum. v^2 is v_x^2 plus v_y^2 plus v_z^2 , which we saw from the vector diagram. From symmetry we know there is no reason to favor the x-direction over the y-direction or the z-direction, and they each get their fair share, a half $k_{\text{Boltzmann}}$ times T . That is the equipartition theory in the simplest form.

Given that, we can then say from the previous viewgraph that the pressure is going to equal mv^2 , you remember that, divided by the volume, and now I have divided that by three, because I am not using v_x here, so it is going to be one-third of v^2 . I can write that as $1/2 mv^2$ times $2N/3V$, where N is the number of particles. Remember, I calculated the force only per particle and we have got a whole box full of N particles.

We end up with this very simple result: $Nk_{\text{Boltzmann}}T$ divided by V . If we express this N in terms of Avagadro's number, we recover $pV = nRT$. So now we have the phenomenological description, $pV = nRT$ (eq. 1), and we have a microscopic model that produces this result. By Einstein's criterion everything is okay: We have got a microscopic model that agrees with the result of the phenomenological description.

DR. HOFER: You skipped over ensemble averaging, really, here, because doesn't the equipartition theorem really refer to ensemble averaging, which is a little bit different from what you said?

DR. GARRETT: Absolutely, yes.

DR. HOFER: It might be worth describing that, unless you want to hurry on past it.

DR. GARRETT: What Professor Hofer says is, again, extraordinarily important. That is why we are here in summer school. I do not have to do 42 viewgraphs. What Tom says turns out, of course, to be the reason that we are not speaking German.

Werner Heisenberg was the top-ranking physicist in Nazi Germany during the Second World War. He calculated the critical mass required for a self-sustained nuclear reaction. He did this by asking how many times would a neutron have to make a collision with a U_{235} and create two more neutrons, in order to get a self-sustained reaction. He realized that there had to be seven collisions before the first neutron got out in order to sustain the reaction.

Therefore, he took the mean-free-path of a neutron in U_{235} and multiplied that by seven to get the radius of the critical mass. He told the army that they needed something like two tons of U_{235} . It could not be done. He made exactly the same mistake that Tom accused me of making, which is the mean free path is the MEAN-free-path: there are shorter paths, there are longer paths.

Actually, the correct calculation, done by Fermi, showed that you needed only 0.7 mean-free-paths, not 7. He was off by one order-of-magnitude in radius and consequently three orders-of-magnitude in the amount of U_{235} that was required. Los Alamos guys, will you bail me out on this?

FACULTY: Don't know nothin', man.

(Laughter)

DR. GARRETT: Anyway, what Tom is pointing out is that is an ensemble average is required. There is a Boltzmann distribution. The Boltzmann distribution is the equivalent of the statement the car makers give you: "Your mileage may vary." Your velocity, as Tom points out, will not necessarily be v_x . As I say, it was an extraordinarily important thing. Werner Hisenberg blew it and Enrico Fermi did not. So I thank Tom for pointing that out.

DR. HOFER: The idea of the ensemble is that you have a bunch of similarly prepared experiments. You have hundreds or thousands of experiments that are all the same, with the same kinds of initial conditions, and you let them go and you average each of these different experiments. So even when you talk about distribution, you are really talking about distribution over similarly prepared situations.

DR. GARRETT: We will save further discussion of the ergotic hypothesis for this evening, when I have a full bottle of cognac.

Anyway, there is a lot of stuff buried in here that I am going over and it is the job of the discussion leaders to nail me. Tom is a discussion leader and he is going to nail me.

Let us go back to the phenomenological model again, since I have treated the microscopic model at such a shallow level. It was discovered fairly late that energy was conserved. Carnot, I believe, wrote his laws of performance about 1822. The reason I bring this up is there was an awful lot that engineers were able to do to make machinery without knowing this law, without really knowing the laws that we hold so dear and think are so fundamental to our existence.

Along the line, a guy named Meyer, who was actually a biologist, I believe, came across this law by doing input-output analysis on animals, measuring the energy they took in and the amount of work they did. Then it was redone by Joule between 1843 and 1849. They came up with this very important result: energy is conserved.

If we have a system and heat comes into the system, it can change its internal energy and it can do work (eq. 3). We have the heat input, which is TdS -- the S is the entropy, which I will

not go to the trouble to define at this point (or at any point) -- is equal to the change in internal energy of the system plus the work that is done by the system. If energy is conserved, this is a conservation law.

(Transparency #8)

Given that result, we can define a quantity that is very useful in talking about heat and sound and materials: that is, heat capacity, or the amount of heat it takes to raise an amount of material by one degree (in this case, 1°C or 1°K), or the change in internal energy with temperature.

(Transparency #7)

From the previous slide you can see that there are two ways that we can do this. We can measure the change in internal energy while we allow the system to do work or when we do not allow the system to do work -- when we clamp that piston.

(Transparency #8)

It is very easy to define the heat capacity from that equation if we hold the piston fixed. The volume does not change. Therefore, the heat capacity at constant volume is just the amount of heat you put in divided by the temperature change (eq. 4).

If we now allow the volume to vary, that is, if we let the pressure be constant and let the piston do work when we add heat, then it takes an additional amount of heat to raise the internal energy, because you are going to raise the internal energy (temperature) and do work.

We can take the ideal gas law and differentiate it (eq. 5). It might as well come out now -- the earlier, the better -- I am a computation cripple. I do not do math. It is a problem. So I have been able to reduce the required amount of mathematics to do science. Dr. Keolian is crippled as badly as I am; there are various people around here who just do not do math.

But there are three things you need to know in math. There are only three mathematical techniques of special value that have great power and they can get you through. One is the most powerful technique in mathematics, which is substitution. The second most powerful technique in mathematics is the Taylor series. The third most powerful technique is chain rule. So I am starting you out with No. 3 here.

The differential pV is $p dV + V dp$ (eq. 5). The differential of R times T , since R is constant, is just $R dT$. So using the most powerful technique, I can substitute this back into the first law of thermodynamics, and I find that dQ is $C_V dT + R dT - V dp$ (eq. 6). In more elite circles this is known as a Legendre transformation, but the name is not important.

Therefore, if I hold p constant rather than V constant, this term disappears. I find that the specific heat at constant pressure is just R greater than the specific heat at constant volume for an ideal gas (eq. 7). That is, this goes to 0, I divide this by dT , I get $C_p = C_V + R$, for one mole of an ideal gas.

Again, this is a phenomenological result based on the first law of thermodynamics, and it tells us that the specific heat at constant pressure minus the specific heat at constant volume is R . This is independent of the constituent of the gas, assuming it satisfies all the assumptions.

What phenomenology does not tell us is the value of C_V , so it is not yet a complete description. We know that $C_P - C_V$ is going to be R , but we do not know what C_V is, so we do not know what C_P is, so we do not have a complete description. We have to revert to a microscopic model, which Professor Bass will do, to answer that question completely.

I went through this, again, just to illustrate that the phenomenological models can give you certain results, the microscopic models can give you certain results, and you have to integrate both of those models to get a complete description.

(Transparency #9)

Because we are talking about acoustics, I want to derive for you the adiabatic equation of state. Adiabatic comes from the Latin *a*, meaning "not," *dia*, meaning "through," and *bainen*, meaning "to go." That is, nothing goes through. The only reason I put this up here is that I am going to assume that the heat capacity of the gas is temperature-independent. I am going to make that one additional assumption.

Therefore, I can talk about any system with a temperature-independent heat capacity as being a polytropic system. The reason I wrote "adiabatic," meaning not through to go, is that I do not have the slightest idea of why a system with a temperature independent heat capacity is called polytropic. I have been looking for a long time. Dr. Swift, you do not know why it is called polytropic?

DR. SWIFT: No.

DR. MAYNARD: "All Gaul is divided into three parts."

DR. GARRETT: These are the guys with the classical education? Larraza was worried that physics was going to become the Latin of the 21st century. Dr. Maynard proves that we have nothing to worry about.

If we make the assumption that heat capacity is not a function of temperature, then it is a polytropic process. What I have done here (eq. 8 and 9) is I have written the generic heat capacity; not at constant pressure, not at constant volume. It is just C . You can do that, you can define the generic heat capacity.

In the case we looked at before, it is $C_V dT$; that is, the heat is C times the change in temperature, $C_V dT + p dV$, or it could be defined as $C_P dT - V dp$ (eq. 9). We can take the ratio of those two quantities and we come up with equation 10.

That is, if I bring this over to that side, I have got $(C - C_P)$ times dT is equal to $-V dp$. I take the ratio of those, I get this result (eq. 10). If I take the V/dV over to this side, it shows that γ prime times dV/V is equal to dp/p . I integrate it. We know that the integral of dx/x is a

natural log plus a constant and, therefore, we end up with the result for a polytropic process: pV raised to this gamma prime power (some constant I am calling gamma prime for reasons that you will see shortly) is equal to some other constant (eq. 11).

If we then make the further assumption that the process is adiabatic, then dQ and dS disappear, no heat goes through, nothing gets in. Since the temperature change is non-zero, dQ must be zero, so C , the generic capacity, must be zero. If the generic heat capacity is zero, then gamma prime is equal to C_p/C_v , which is what we know as gamma, the polytropic coefficient. We get pV to the (gamma) power is a constant (eq. 12).

So this is merely a consequence of the fact that heat capacity is not a function of temperature. If, on the other hand, if we let $dT = 0$, then the generic heat capacity must be infinite for any heat to take place (because if it is isothermal $dT = 0$ and dQ is some value, then this must be infinite).

We go back up there and we see the ratio is minus infinity over minus infinity, which is just one, and we get pV raised to the first power = a constant. So we have recaptured the ideal gas equation-of-state: $pV = nRT$. If T is a constant, then pV is equal to a constant, because T is a constant, then RT is a constant.

There is a complete range of equations of state that run from the adiabatic case to the isothermal case and they are characterized by this gamma prime going somewhere between C_p/C_v and 1. A phenomenological result.

(Transparency #10)

Now that we have that well in hand, we can go on to the motion of ideal gases. As I said before, for a homogeneous, isotropic, single-component gas, five variables provide a complete description of the system. I am going to choose p , S , and v : pressure, entropy, and velocity. Why do I choose entropy? Because I am going to talk about adiabatic propagation and, therefore, dS is zero and we automatically contract from five variables to four variables. That is why I chose S instead of T as my thermal variable.

Now it is incumbent on me to write down four equations to close the description. The first conservation equation will be conservation of mass (eq. 10): $d\rho/dt + \nabla \cdot (\rho \mathbf{v}) = 0$. You are all familiar with this. We all agree that mass is neither created nor destroyed.

I want you to appreciate the form of a conservation equation. This will be showing up again. If you look in your lecture notes, almost everybody writes this down at some point or another. The form is important.

The time derivative of a density (in this case, the mass density) plus the divergence of the associated flux density (in this case, mass flux density), is zero. So it is a time derivative of a density and the divergence of a flux. Here the flux is the mass flux; I am writing it as \mathbf{j} , which is $\rho \mathbf{v}$.

That is one equation, because it is scalar since ρ is a scalar. We can write down momentum conservation (eq. 11). That is three equations. Again, it is the time derivative of the density (in this case it is the momentum density, ρv) plus the divergence, which I have written here in tensor form. I assume everybody is familiar with the Einstein summation convention, right? If any index is repeated, we assume that there is a summation over that index?

It is the time rate of change of a momentum density and it is the divergence of a flux, where this flux is the momentum flux (eq. 15). It is the pressure times the Kronecker delta ($\delta_{ik} = 1$, if the indices are identical, or 0, otherwise.), plus $\rho v_i v_k$.

Since I want to stick with vectors, I can rewrite this conservation equation (eq. 16), because I can take this term over here in the divergence and write it in a more familiar form as Newton's second law, or Euler's equation. You can see that this is equivalent to that; this $v \cdot \nabla$ term is just this term with that stress tensor.

Now I have got four equations and I have got five unknowns: I have got density and pressure and velocity. So I am back to five unknowns. Again, our phenomenological description, our thermodynamic description, will bail us out.

(Transparency #11)

Because we know p for the fluid, pressure is a function of density and temperature, density is a function of pressure and entropy, and that any two variables are a complete description of the system -- that is what phenomenology told us in the ideal gas law -- therefore, I can use the No. 2 most important theorem in mathematics, Taylor expansion, and expand the density about its equilibrium value.

I can say the density is the mean density, or the ambient density, plus $\partial\rho/\partial p$ at constant S times dp plus $\partial\rho/\partial S$ at constant p times dS -- two variables are enough to describe the system -- and I can continue this Taylor series: $1/2!$ times $\partial^2\rho/\partial p^2$ at constant S times dp^2 plus the same thing for this, plus a cross derivative, *ad infinitum*.

For Mark Hamilton's talk on nonlinear acoustics this will be important. This is where he introduces a thing called the B/A coefficient. I think Dr. Hamilton does stick with adiabatic processes, so all the dS 's go away, but this coefficient here is your B . And this coefficient here is your C . I have truncated it already.

But we are going to stay at the linear level. We are going to say that it is an adiabatic process, as I assumed in the beginning, so all the dS terms get chopped. And we are going to look at only small deviations from equilibrium and, therefore, dp , the change in density, which is ρ minus this ambient value, will truncate the series (eq. 18), and it will be $\partial\rho/\partial p$ at constant S , times dp .

All we are doing is a Taylor series expansion (eq. 17), with the knowledge that two variables form a complete description; that is, I did not have to take this dp , dS , dT , blah, blah, so it is only two variables and it is complete.

(Transparency #12)

Acoustics: Finally! The science of the subtle. We expand the variables about equilibrium. We say that the pressure is some ambient pressure plus some first-order deviation from equilibrium plus some second-order deviation from equilibrium, et cetera. The same is true of the velocity.

P_2 is understood to be proportional to the first-order quantity squared. I am going to talk about systems that do not have any flow superimposed on them; that is, v_0 . There is no dc flow, there is no Doppler shift.

If we linearize the equations and retain only the first-order terms, we can assume a solution in terms of a complex exponential. This turns out to be very useful. The pressure will be the real part of this expression and the velocity will be the real part of this expression, where it is understood that v itself may be complex; that is, we can put a phase shift in.

There is a phase speed which is given by the ratio of ω over k or, more familiarly, the speed is frequency times the wavelength.

Remember, when we did the Taylor series expansion of the equation-of-state, the change in density is ρ_1 , a first-order deviation from equilibrium is $\partial\rho/\partial p$. This is at constant times p_1 (eq. 19). I am going to call $\partial\rho/\partial p$ at constant S , a^2 . It is a constant which turns out it is a positive constant; otherwise the universe collapses. So a^2 is just fine.

If we now take the momentum equation and this harmonic solution assumption, we can substitute that in and convert the differential equations to algebraic equations. That is the whole purpose of making this assumption. When I take the $\partial\rho_1/\partial t$ from the mass conservation equation, the derivative with respect to time gives me a $j\omega$. The ρ_1 I want to express as p_1 , so it is just p_1/a^2 , due to the equation-of-state.

The same thing is true with the divergence. It is $e^{j(\omega t - kx)}$, so I pull out a minus k ; dot it into v , and it is equal to zero (eq. 20). So all I have done is transform this linearly and algebraically. I do the same thing for Euler's equation (eq. 21), $\partial/\partial t$ of ρv gives you this $j\omega$ and the gradient of the pressure gives me the $k = 2\pi/\lambda$ times the pressure.

You can see I have a system of two coupled linear equations, and for those equations to have a consistent solution, the determinant of their coefficients, that is, $j\omega$ over a^2 times $j\omega$ times ρ_0 , is going to cancel out, minus the other cross-term in the determinant. Those must be equal to zero. For a linear set of equations to have a nontrivial solution, the determinant of their coefficient must vanish. Now I said that almost like I knew some math, but don't be fooled.

If we take that, we find we get j times j is a minus, ω times ω is ω squared, we get a rho naught coming up this way, we get a minus times a minus, which is a plus, times a j times a j , which is a minus -- that is the other diagonal, which is a minus, which is a plus -- and you end up with ω/k is the square root of p/ρ , or this constant a .

DR. ATCHLEY: Would you repeat that?

DR. GARRETT: No.

(Laughter)

It is in the notes -- probably wrong, but it is there.

DR. HAMILTON: Do you want to say a word about why you would assume the process is adiabatic and not isothermal?

DR. GARRETT: I am going to say in about three to five viewgraphs why it is adiabatic, absolutely. I am going to harp on that.

DR. HAMILTON: After all, Newton made that mistake.

DR. GARRETT: Yes. But he tried to cover up. That is what was important. He really set the stage and he covered up. It took a hundred years to bust him.

(Transparency # 13)

What I would like to point out at this point is that some of you might say I have used mass conservation, I have used momentum conservation, but I have not used energy conservation and, therefore, I have another conservation law at my disposal and the system is, therefore, over-determined.

That is not true. Energy conservation, as I show up here, is a consequence of mass conservation and momentum conservation. That is, if we multiply the mass conservation equation, which is eq. 13, in your notes, by p_1 , you get $p_1 \bullet \nabla v$ is equal to p_1 over ρa^2 times $\partial p / \partial T$ (eq. 23). If we dot multiply v_1 into the Euler equation, we have $v \bullet \nabla p$ is equal to a minus $v \bullet \partial v / \partial t$ (eq. 24).

Again, from chain rule, it is very easy to see that this term, p_1 times $p_1 dt$, is really $(1/2) p_1^2 dt$. The one thing I can do in calculus is take the integral and the derivative of polynomials. I am good at that. I know you take one away from the exponent and you add one. I've got that.

So this can be written as $(1/2) \partial v^2 / \partial t$, and this can be written as $(1/2) \partial p^2 / \partial T$. Again, from the chain rule, that the divergence of pV is $v \bullet \nabla p$ plus $p \nabla \bullet v$. When we put that all together, lo and behold, we find $\partial / \partial t$ of this quantity plus the divergence of this quantity is equal to zero.

It does not take a great deal of imagination to identify the first term in equation 25 as the kinetic energy density. So we have got the time derivative of a density. Even if we do not know what the second term is, we know that the first term is kinetic energy, so the second term must be potential energy.

The time derivative of the energy density is the divergence of the intensity, which is the flux. I bring this out for a couple of reasons. One is that Greg Swift is going to work quite a bit with acoustical energy transport in fluids. The other is that conservation of energy is not an independent conservation law. It is contained within the other two conservation laws.

DR. MIGLIORI: Having not thought about this at all myself, I certainly do not need to conserve momentum at a boundary.

DR. GARRETT: Absolutely. There are no boundaries.

DR. MIGLIORI: So we are going to have to be very careful about statements like that when we get near boundaries.

DR. GARRETT: Absolutely true. What Dr. Migliori has pointed out is that I have made certain assumptions that are reflected -- I think this was what caught his attention -- in the fact that energy is not necessarily conserved; it can be dissipated. But I have left thermal conduction out of the system, I have left viscosity out of the system --

DR. MIGLIORI: Or momentum.

DR. GARRETT: Viscosity shows up in later viewgraphs.

DR. MIGLIORI: Really, the energy conservation rule is the fundamental one and the fact that the others -- well, anyway.

DR. GARRETT: Okay. Professor Migliori is in what I call the Greek-Babylonian kind of

--

DR. MIGLIORI: "Babble" would be fine. You could just stop with "babble."

(Laughter)

DR. GARRETT: No, this is valid. There are two ways of doing mathematics, of doing logic. One is the Greek way, in which there are theorems that have a great deal of weight. On top of these theorems you can build up all these corollaries and lemmas and the other things. But there are certain things that are more important than other things, like two parallel lines never meet. That is a theorem, the Euclidian view, the Greek view of mathematics.

The Arabs, on the other hand, the Babylonians, had a different view. They invented algebra. They did not invent geometry. To them there was no hierarchy. Algebra allowed you to get from any one statement to any other statement by this technique we call algebra and, therefore, no statement was more valuable. That is, there was no hierarchy, there were not theorems and corollaries and lemmas.

What Dr. Migliori has brought up is the question of what is the fundamental expression, and I will give you the Babylonian answer. Do you speak Arabic? Can you say "It don't make no never mind?"

DR. MIGLIORI: Maybe I will bug you some other time about this.

DR. GARRETT: No, it is perfectly fine, but is this more fundamental than momentum conservation?

DR. MIGLIORI: The thing is, if you put in scattering centers somehow, like stuff, and you look at how the temperature increases if you have flow and you are scattering off of hard objects that do not move, it is very nice to have conservation of energy right there at your hip.

DR. GARRETT: I am going to start introducing boundaries. It is on the agenda. Whether I get to it or not is irrelevant. But you are absolutely right. The assumption I have made here is a nonthermal conducting, nonviscous fluid. Therefore, energy is conserved. If it is not conserved, I think it is worth pointing out, as long as Dr. Migliori brought it up, that there could be source terms in equation 25.

It is the same thing with the mass equation. In fact, I was forced, after 13 years of successfully avoiding teaching the theory of radiation in acoustics. You can very simply introduce dipole radiation by taking the momentum equation and allowing external forces to act on the fluid. So momentum is no longer conserved. There is a source term in force. There is a source term in mass that produces monopole radiation. Things can spew and suck fluid in the neighborhood, and you get radiation out of that.

But what I am doing is this thing where there are no sources or sinks of mass or energy. But you are right. It gets much more interesting and much more powerful if we include those things. And good, I cover myself here, right? This is a nondissipative result. Sources could include viscosity and thermal conduction.

Another thing I wanted to point out is that the energies are quadratic in the linear fields. That is, we did linear acoustics. We linearized the system. But when we came out to look at the energies, everything is nonlinear. Everything is quadratic; the product of two linears is bilinear terms or straight quadratic terms.

(Transparency #14)

We can calculate the speed of sound -- this should be under isothermal conditions, as Mark Hamilton pointed out, the Newtonian result. We know that p is equal to ρ times RT/M . If T is a constant, then the derivative, $\partial p / \partial \rho$, at constant temperature is a^2_N (eq. 26).

I can take this derivative, the derivative of p with respect to ρ , which makes that a 1, and it is RT/M . This is known as the Newtonian sound speed squared, RT/M . If you look in the *Principia*, 2nd Edition, 1713, that is the speed of sound. It is about 980 feet/second -- he was British, we use feet/second.

The experimental value was more like 1100 feet/second. So Newton hemmed and hawed and in various editions he put in various other things, but the fact is, he got it wrong, because he assumed that the propagation was isothermal. It was Laplace, about a hundred years later, who

realized, as I will show in subsequent viewgraphs, that the propagation is, in fact, adiabatic in gases.

If it is adiabatic, you can write pV^γ as p^γ , because density and volume for a fixed number of moles of gas are inverses. We can take the natural log and differentiate this expression. If we take the log of this and if that is then equal to minus gamma times the log of that and that is equal to the log of a constant, we take the derivative, the log of a constant disappears. We get dp over p naught is equal to γ over ρ naught times $d\rho$ and, therefore, $dp/d\rho$ is γp over ρ . That is what I have shown here (eq. 27).

This sound speed, a_s , is γp naught over ρ naught. From the ideal gas law we know that p naught over ρ naught -- remember, I did the intensive form of the ideal gas law -- is just RT/M . So the sound speed for adiabatic propagation is $\gamma RT/M$ in ideal gas (eq. 28) and, for isothermal propagation it is RT/M -- which, again, we expect from our polytropic analysis, where, for the isothermal case, gamma was equal to one. In the adiabatic case gamma was equal to the ratio of specific heats.

So there is this factor in air of 1.4 that stiffens the gas adiabatically.

(Transparency #15)

If, in fact, the adiabatic compressions take place, then the temperature cannot remain constant: It is either adiabatic or isothermal or somewhere in-between. If we assume it is adiabatic, there will be temperature changes. Again, taking the adiabatic gas laws and substituting (the No. 1 most powerful technique in mathematics) for ρ to the minus gamma, which comes in here, we get the expression in equation 29.

If we collect the terms, we get $p^{(1-\gamma)}$ times T^γ is equal to some other constant, constant double prime (eq. 30). Again, we take the log and differentiate, we get gamma minus one times p_1/p_0 -- that is, the deviation of pressure over the ambient pressure -- is equal to minus gamma times the deviation in temperature over the ambient temperature (eq. 31).

We are left with an expression that has kept Dr. Migliori and Dr. Hofler and Dr. Swift and me and several other people employed doing thermoacoustics. When you have a sound wave, when there is a pressure change relative to ambient, there is also a temperature change. The size of that temperature change is related to this polytropic coefficient. That is why understanding gamma is so important.

In the obvious limit, where gamma equals one, if you have an isothermal process, this goes away, which is correct -- you always want to check that. In the case where gamma is greater than one, then a pressure change is accompanied by a temperature change. This is not normally emphasized in most acoustics classes, because for normal speech, say, 74 dB, which is 0.1 Pa RMS, the temperature swing, the RMS variation in temperature due to my voice at 74 dB is 83 μ K -- it will not cook a chicken. It is not a large amount of temperature change.

On the other hand, for a thermoacoustic refrigerator -- the one on which I happen to have the data closest to hand is a fridge that we put out on a naval vessel, so it is called SETAC, shipboard electronic thermoacoustic cooler -- the pressure swings there were 65 kPa. These are large amounts of sound. It is about a 190 dB sound pressure level. One hundred thirty dB will damage your hearing, 140 dB will damage your hearing and hurt you. One hundred sixty-five and above will light your hair. The sound-wave amplitude is large enough that the viscous interaction between the gas and your hair creates enough friction to make hair catch fire. So we are talking about *beaucoup* dB.

We are not talking about stuff you normally run into, but this is the RMS pressure in a typical thermoacoustic refrigerator.

If we take a look at the ambient pressure, which in the SETAC case was 20 atm, gamma for the helium-argon mixture is 5/3, the ambient temperature is taken to be 20° C, the temperature of swing now, is 3.6° K. The peak-to-peak pressure is almost 20° F. That is the difference between a comfortable room and Mississippi.

(Laughter)

This effect can be very large. In fact, that effect can then be cascaded. That technique was the regenerator by Rev. Stirling and called the stack by Swift, Migliori, Wheatley -- or Swift, Hofler, or Wheatley, Swift, and Migliori, or some combination like that.

DR. HAMILTON: The argument you seem to be making is that it is an amplitude-dependent effect; the louder the sound, the more likely it might be to be isothermal as opposed to adiabatic.

DR. GARRETT: No, no, absolutely not. What I am saying is that the larger the sound amplitude, the larger the temperature swing. For modest sounds in everyday experience in air, that temperature swing is entirely negligible. But if you crank the acoustics up high enough, you get substantial temperature swings that accompany the sound wave because it is adiabatic. If it were isothermal, there would be no temperature swing.

DR. RASPET: It cannot be negligible since it changes the speed of sound by a factor of the square root of 1.4.

DR. HAMILTON: But the transition from adiabatic to isothermal is a frequency-dependent effect.

DR. GARRETT: I am going to show you that. Absolutely. It is very important. Dr. Hamilton is very careful, because he knows this is important for him and he does not want me to mislead you yet.

(Transparency #16)

What are the problems so far? The specific heat is undetermined, because we have not looked at the microscopic (atomistic) picture, and we do not know why adiabatic is correct and isothermal is wrong. That is what we are about to get to.

Let us summarize where we have gone so far, just to get everybody recalibrated. I introduced the phenomenological approach, which was based on the following question: How many variables are required to form a complete description of the system? At that point we closed the description by applying the appropriate number of conservation laws. These may be energy conservation, mass conservation, Galilean invariance, equation-of-state, what-have-you.

We then applied the phenomenological in the case of an ideal gas. We derived an adiabatic equation of state, making the assumption that the heat capacity is independent of temperature. We found from energy conservation that the difference between the specific heats for an ideal gas was equal to R (eq. 7). The ratio of the specific heats we have named as a constant, because it was important -- γ , the polytropic coefficient (eq. 12).

Under adiabatic conditions, pV^γ was equal to a constant (eq. 12). The sound speed squared is $\gamma RT/M$ (eq. 28). The temperature swing associated with that adiabatic compression and expansion is given by γ minus one over γ times the pressure ratio times the ambient temperature (eq. 32).

As Dr. Hamilton points out, this description leaves a couple of things to be desired. One is I have not told you how to relate γ to the microscopics of the particular gas, although I think Professor Bass is going to do that quite carefully. One then must go back into the microscopic domain and introduce quantum mechanical concepts regarding degrees of freedom and energy gaps between rotational modes and vibrational modes and other things, which are much more suited to Professor Bass' lecture than they are to mine.

The other question is the one Dr. Hamilton is serious about: "why is sound adiabatic and not isothermal?" Why was Newton wrong? Before I answer that, let me talk about a couple of applications of this rather important result: The speed of sound in an ideal gas under adiabatic conditions is $\gamma RT/M$.

If you satisfy those conditions exactly, that equation is exact. In fact, as Dr. Swift pointed out, this equation has been used to a measure fundamental constant. Because of the kinetic theory, as Professor Bass will point out, for inert gases γ is exactly $5/3$. At sufficiently low pressures that result is true.

Therefore, if we can measure the temperature, and we can make an accurate measure of the speed of sound, and we know that we have got, say, helium in there, which has an atomic mass of 4.0026 -- and another digit I do not remember -- then obtain R . In fact, several years ago that was done by Mike Moldover, Mo Greenspan, Jim Mahl, and their cronies at the Bureau of Standards.

They built a spherical resonator, which I will show you (not their resonator, but another spherical resonator). The reason it is spherical is that you have a higher Q, since in a radial mode there is no scrubbing, as Dr. Migliori was talking about -- we just had thermal conduction losses, which we calculate a few viewgraphs down the line.

But they were able to make an accurate measurement of resonance frequency and an accurate measurement of absolute temperature. In fact, it was so successful that these results made the front page of Japanese newspapers. For American newspapers, it did not involve people getting AIDS, or actors, or anything, so it never even made it. It did not make it to the back pages; it did not make it at all.

But it was an important result, because this is the second worst known constant in physics, and they improved the measurement acoustically by a factor of five -- it was not as if they improved it by a few percent. They improved it by 500 percent; that is, they reduced the uncertainty by 500 percent.

So our expression for the speed of sound is useful for making fundamental measurements and, in fact, it was so useful that the guys there are going to use it to redefine some of the fixed temperature points. They are going to use it for an ideal gas absolute thermometry.

(Transparency #17)

Another application of this in thermometry is far less precise but extraordinarily useful, and that is acoustic pyrometry. This is a diagram of a basically time-of-flight apparatus that is installed in large gas flues in power companies. If you look out there when it clears up, you will see two big stacks for PG&E -- they have a big power-generating plant at Moss Landing.

Besides the electricity, it generates hundreds of millions of dollars in revenue. Therefore, if you can improve its efficiency by a percent, it saves the company a million bucks. A small improvement in efficiency can have a very large result in your bottom line and, depending on the cost of fuel and the capital costs, it can be the difference between a Dilbert cubicle and a carpeted office with a view.

What has been done recently is to develop a system that sits in the stack that has basically two whistles and two microphones. They do acoustic tomography in the stack. That is, this whistle uses plant compressed air, which is pressurized gas (pressurized air exists in almost every major industrial plant on the planet). It has a whistle and a whistle, sort of a foghorn whistle, and a microphone at that whistle.

It is very nice to use a whistle in this environment, because it is a stack, it is loaded with crud. You put a microphone on there, or a speaker on there, and it would just get covered with crud after a sufficient amount of time. So this whistle is a self-cleaning acoustical source. It goes "bloop" and blows all the crap out and this guy and this guy listen. Then it is this guy's

turn, and then this guy whistles and this guy listen. They do time-of-flight tomography and this is what you get.

You find out that in your burner the hot part has moved away from the center. These are equal temperature contours (isotherms) and these are big temperatures. You put a thermocouple in here and it becomes liquid, it becomes vapor at these temperatures. But there is hot stack gas going through there and the time of flight in both directions, averaged, gives you the speed of sound in the gas constituents (I will say more about that later).

By passing different paths through here, you get a variation in the speed of sound and you can plot out the corresponding temperature and, therefore, you can open this damper and get that thing closer to the middle. You can do that with acoustics. You can sell these things.

This simple result, if you go to the trouble of industrializing it, makes a very useful and very profitable device.

MR. LAVRENTYEV: In the previous example there was an equation where you can control temperature. You can define what temperature. You said that when you have high vibrations, temperature should be changed.

DR. GARRETT: Right, but it changes at a high frequency. It changes at a frequency so that as you pass gas through over many cycles you get an averaging. There is still an ambient temperature and it is a little bit hotter and a little bit colder in the crests and the troughs, but you are pushing a lot of that through.

It takes the average itself, physically takes the average, which is important, it is very important.

(Transparency #18)

One of the things I have not pointed out is that we can also use this effect -- $\gamma RT/M$ -- to determine M . If we have the molecular weight of a mixture, say, of two components (component one and component two), then the mean molecular weight is the average of those, the concentration X times M_1 plus one minus the concentration times M_2 (eq. 33).

We can do a similar weighting for the polytropic coefficient, but we cannot do it this way. We cannot say X times γ_1 plus 1 minus X times γ_2 , because these are intensive variables and these are extensive variables. To do it correctly, you have to have X times C_p of component one, 1 minus $X(C_p)$ of component 2 over -- the same thing for C_v to give you γ (eq. 34).

But it turns out that this is a pretty good approximation, as I show down here. Therefore, the speed of sound in a mixture over the speed of sound in the pure component No. 1 is given roughly -- that is, approximately -- by this expression in equation 35 which simply involves X times the relative difference in the γ s and X times the relative difference in the masses.

The resonance frequency in a resonator is proportional to the sound speed, therefore, the change in the sound speed over the sound speed of the first component is this combination of quantities times a change in concentration over two. I have defined this whole piece of junk as beta, the sensitivity factor for the dependence of resonance frequency on the gas concentration (eq. 36).

What I plotted below equation 36 is the line, M_2 minus M_1 over $2M_1$, which is the result that you would get for beta if you forget all about gamma. Shown on that graph are roughly 20 different gas constituents where I have done the calculation accurately. You can see they lie very close to that line. So for purposes of discussion at the moment, we will not worry about the gammas.

(Transparency #19)

This turns out to be one of the best methods for measuring contaminants in gases. It was so successful that it was in use commercially as early as 1884. In Germany there was a device called the *Schlagwetter-pfeife*, the wet gas whistle. In mines, in Germany, they pump air through the mines so the guys digging inside did not die of suffocation. Instead of just exhausting the air into the atmosphere they would blow it through a whistle.

If there was methane in the mine, or hydrogen in the mine, which can cause explosions, the pitch of that whistle would go up. The whistle has constant length, so the wavelength was fixed. The speed of sound went up, the frequency went up. That worked for all but the tone-deaf miners. If you were tone deaf, you could not hear the change in pitch, you would die, because you would not leave the mine.

To save the lives of the non-musical miners, they used two whistles. They took one with house air and one with mine air and you could hear the beats, so even the tone-deaf miners did not croak. If you heard this whistle going -- [undulating whistle pitch] --you knew it was time to get out of the mine, because there was hydrogen or methane.

Very simple, very accurate, very sensitive. The temperature, of course, did not matter, because these two whistles were side by side, nice metal whistles, and the temperature was uniform.

So it was only the gas constituent, the mean molecular weight, that mattered. That would tell you that something was awry in the mine. Hydrogen is lighter in weight than air and methane is lighter in weight. The pitch would always go up when there was a danger of explosion.

The acoustical technique actually fell out of favor at about the turn-of-the-century, because it was cheaper to measure voltage than frequency. From the turn of the century on out the measurement instrument of choice was the Wheatstone bridge. They had very sensitive galvanometers and you balanced the bridge. During that time, from 1888 to roughly now, people

used hot wires to measure the thermal conductivity of the gas mixture, not the sound speed of the gas.

You would take two hot wires, one in the pure gas and one in the unknown, and since the thermal conductivity was also proportional to the square root of the mass, it would unbalance the bridge. But it is a very difficult measurement.

Nowadays that is not true. This is a wristwatch that cost \$20.00 (I do have a nicer one, but this is the one I use). It is a frequency standard that is accurate to 10 ppm; it is good to a part in 10^5 for only \$20.00. A voltmeter that is accurate to a part in 10^5 costs well over a thousand dollars.

So with the evolution of technology the Wheatstone bridge does not have the highest precision-to-cost ratio anymore; the highest precision-to-cost ratio is now frequency or time measurement. Since sound speed in a resonator modifies the frequency, there has been interest in developing acoustic as opposed to thermal conductivity measurement techniques.

This work was initiated in the early 1970s by Ray Kagiwada, who is a colleague of -- I guess, Moises (Levy), he was around when you were there, right, but he was after you were there, right, Dick (Stern)? There are a lot of guys from that group at this summer school.

They were working in low-temperature physics. They were looking at helium-3 and helium-4 mixtures and they wrote a paper on the measurement of isotopic ratios of helium-3 and helium-4. This was published in *The Journal of Low Temperature Physics*, because these were the only guys who cared about helium-3 and helium-4 mixtures, unless you were building a nuclear weapon.

Jim Frazier was at Los Alamos, where nuclear weapons are developed, ergo, he was also interested. He was another Rudnick student of the Stern/Levy era. In 1972, he actually built one of these acoustic resonators for gas analysis. I was visiting Greg Swift at Los Alamos, John Wheatley's group out there, and I actually found Frazier's original resonator for making these measurements.

Subsequently, there were all kinds of whacko applications of this, one of which was a device that I am going to show you that was produced by Robert Keolian, who is in the back, known as the "brass balls." It was a pair of spherical resonators cut out of a single enormous chunk of brass, so the two spherical cavities were in excellent thermal contact.

The "balls" were a differential analyzer; that is, you had a pure species in one of the balls and you had the unknown species in the other ball. I will show you, it is a beautiful instrument. It took a lot of work and ate up a lot of time. It had a terrible narcissistic effect on Robert: He fell in love with the thing to the exclusion of all kinds of other things. They had a short tumultuous romance and it was over. I have pictures, as you are about to see.

When I was at Berkeley, we had a problem because the helium in Berkeley was recovered and re-liquefied. The chemistry department had a helium liquefier. You recovered the helium out of your low-temperature Dewar and blew it back into this liquefier. If it had too much air in it, it would shut the liquefier down, it would plug up the liquefier.

So Dr. Swift and I put together a gas analyzer that was not nearly as sophisticated as the one that Robert Keolian put together. It consisted of a very simple cylindrical resonator with two electric transducers on it. Dr. Swift took about 10 meters of copper wire which he used as a thermometer by winding it around the resonator. The mike and speaker were wired to be self-resonant.

You would take the microphone output, amplify it, shoot it back into the speaker, and it would sing at its resonant frequency. There is a little inductor in there to keep it from mode hopping to the higher modes. It was a very simple device.

The frequency would get converted to a voltage with a \$2.00 tachometer chip, LM-2907 or 2917, something like that, so that the frequency would become a voltage. The resistance of the copper wire would be converted to a voltage. Those two would be summed electronically to compensate for the temperature difference, which is important, because the helium would come out at very low temperatures as it passed back into the recovery system when it boiled off from the liquid helium at 4 °K.

Then there would be a comparator with a meter that would read out directly in percent of air and then an alarm would go off if we had to shut it down and take it off the chemistry recovery line. There is a whole span of gas analyzers from the simplest to the most baroque. I will show you the two examples of that shortly.

The limitations on this technology are really not very severe. As Dr. Hamilton will point out, if you measure this B/A coefficient, which tells you the speed of sound dependence on pressure, for air the pressure dependence, that is, $(1/F) \partial F / \partial P$, is only 38 ppm/psi. So even the most horrendous changes in atmospheric pressure do not significantly affect the sound speed.

On the other hand, a humidity change of ± 20 percent will give you about an ± 800 ppm change in the frequency. If it is important, again, just as we had a thermometer here, you might have to have some humidity compensation.

DR. SWIFT: In the chimney or industrial burner you were talking about, how did they deal with the humidity problem? They would have to know what effect the gas constituent would have. If the sound speed is T/M , if you want a negative thermometer, you have got no M .

DR. MIGLIORI: It is in a stack, right?

DR. SWIFT: I was not sure if you started out in a chimney or in a burner.

DR. GARRETT: It is in the chimney part.

DR. MIGLIORI: So if they are making money on burning stuff, if they are going to burn it, right, with gas concentrations? Maybe they do not care about 100° here and there. The hydrocarbon concentrations of fuel vary wildly. Maybe they just do not care.

DR. GARRETT: They get the mean sound speed, right? If you broadcast from this guy to this guy and measure the time, and broadcast from this guy to this guy and measure the time, and take the average, you get the sound speed.

DR. MIGLIORI: They know the sound speed. The problem is how does that relate to the temperature if the density of the gas varies?

DR. GARRETT: The assumption is that in the stack the constituents are well mixed, that there is not variation from place to place. There are not more hydrocarbons here and more something else here. And that is an assumption. But apparently it is a fairly good assumption.

(Slide)

I did want to show you this device. I do not have a picture of the National Bureau of Standards (sic) spheres, where they redetermine the gas constant, but I have pictures of these spheres that were built by Robert Keolian. He had fallen in love with it by this time. He had built it a little stand and it had knurled handles so you could carry it around, and dates -- he would take it out to dinner and put it on the table and take it to the movies.

(Laughter)

This is the center section, of which I will show you an open version. These are the end caps. In each end cap there were two transducers. This transducer is called a picophone -- I will show it to you, I have a separate picture of that. Along the wall -- this is a preamp -- there was a condenser microphone.

(Slide)

Here the sphere is open. You can see the inner sphere. This is the fill line. This is the picophone (you cannot see it well, I have another picture of it). It is a polyvinylidene fluoride (PVDF) loudspeaker at the center, so that you drive only the radial modes. There you can see, in poor focus, the face of the microphone.

DR. MIGLIORI: You lost me. It is in the center of the 3-D volume?

DR. GARRETT: Yes, it is at the center of the 3-D volume. It is coming in here on an angle. When this thing closes up, it is right there.

DR. MIGLIORI: It is on a little stalk there?

DR. GARRETT: Yes, you will see it. I think it is in one of the next two slides.

(Slide)

There is the source, on this thin stalk. There is the PVDF rolled in a cylinder. It is a monopole source. PVDF is a piezoelectric polymer. When you apply a voltage between the

inside and the outside surface it gets longer. Its strong coupling coefficient is a $D_{3,1}$. So if you apply a voltage across this way, it grows that way.

So if you wrap it into a cylinder, it pulsates. This is one of the things Dr. Keolian found attractive about this device. This is the electret microphone, which you can see is just a charged piece of thin ($12\text{ }\mu\text{m}$) aluminized teflon sheet. Behind it is the electronic circuitry.

(Slide)

This is a close up of the picophone. It is really nice. This is conductive epoxy on here to make the shield. That whole thing is probably not more than $1/8$ of an inch in diameter.

(Slide)

There is Dr. Keolian 22 years ago, or thereabouts, attempting to mate with this device.

(Slide)

This is the device that Dr. Swift and I built at Berkeley. I showed you the block diagram of it earlier. There is the microphone and the speaker -- here -- the self-resonant circuit. We would calibrate this by changing the frequency on the electronics and it would read out until it got to the point where the chemistry department would take us off line and it would set off that alarm.

This line here is just the connection to the copper wire -- I can see it, because I am right up here -- but it is wrapped around that resonator body. The gas would come along that tube, the PVC tube, up there, back to the chemistry department. And there is a shut-off valve, which you cannot see in this picture.

I think that with time acoustic gas analysis will supplant the thermocoupled-type gas analyzers. They are both good for binary or pseudo-binary gas mixtures. They are not specific; they measure only the mean molecular weight. When you want specificity, people would take what is called a gas chromatographic column; that is, you take a long thin capillary and you push the unknown gas (you gassify the unknown) through this capillary and the big molecules will be slower, their viscosity would be higher and it would take longer for them to get out.

As the helium gas made it through, the lighter constituents would come through and change the thermal conductivity. Then, after a certain amount of time the heavier ones would come through. So even in these analyzers that are species-specific, like the gas chromatograph, at the tail end it is just a thermal conductivity gauge.

I thought that since we had gone through the ideal gas law, the adiabatic equation of state, and the speed of sound, it would be nice to see a couple of applications that involve either acoustic pyrometry or acoustic gas analysis.

I would like to raise the stakes a little bit and introduce the concept of irreversibility. As Mark Hamilton points out, I have not addressed this problem of whether sound is adiabatic or isothermal in any quantitative way.

(Transparency #20)

To do that, we have to understand thermal conduction and we have to understand irreversible processes. What I would like to do in the remaining hour is develop a dissipative phenomenology; that is, take a look not at the reversible equations -- the time-reversal invariant equations -- but now add dissipative processes that generate entropy and, therefore, are irreversible.

The way I would like to do this is to introduce the concepts of thermal conduction and viscosity, which are going to play a role in Professor Bass' lecture on attenuation processes in molecular acoustics.

(Transparency #21)

We start out, again, from the simplest point. For dissipative processes, the model system we want to think about is power dissipation in electrical resistors, because we are all familiar with Ohm's law: The current depends on the voltage difference and the resistance.

The equivalent of Ohm's law for the transport of heat is Newton's law of cooling (eq. 40): The heat flux is proportional to the temperature gradient and the proportionality constant is called the thermal conductivity. (I have shown the dimensions up here.)

So we can get, basically, the electrical resistor result by taking a look at a rectangular slab of material with cross-section A and with length, L , that has an amount of heat, Q , passing through it, as opposed to heat flux, which is Q per unit area. So $Q \cdot A$ is Q .

Therefore, we multiply both sides by area and we get the temperature gradient, which is linear in this case, T_{hot} minus T_{cold} divided by the length. We end up with an expression that looks exactly like Ohm's law (eq. 41): The current is proportional to the change in the voltage and, in this case, the change in the temperature, and that proportionality coefficient is something that is the thermal conductance, just as this is controlled by the electrical conductance. The longer you make the slab, the higher the resistance; the larger in cross-sectional area, the lower the resistance.

You can think of this as saying that if you take two of them and put them in parallel, you double the area and it lowers the resistance. We have gone from a scalar equation to a vector equation, but basically it is Ohm's law.

Similarly, with viscous effects, we have Ohm's law, but it is no longer a scalar equation, it is not even a vector equation; it is a tensor equation (eq. 42). But it is still Ohm's law. The force-per-unit area in the x -direction due to gradients in the y -direction. That is, if I have a plate like this and I have got fluid flowing over it, the fluid at the surface of the plate is stuck to the plate, the fluid above the plate is moving, so there is a gradient in the y -direction that puts a drag force on this plate in the x -direction.

So the y -direction means the direction that is normal to the surface. Then x is the direction of the flow, so the gradient is $\partial v_x / \partial y$. The constant here is known as the shear viscosity (its dimensions are shown).

(Transparency #22)

As I am sure you as acousticians at some point have derived the wave equation on a string or in a thin bar -- I am assuming that everybody in this room has had an acoustics course at the level of Kinsler and Frey. May I see the hands of people who have not had an acoustics class?

(Members of the faculty and the discussion leaders raised their hands.)

(Laughter)

Okay, if we include those of you who have taught an acoustics course at the Kinsler and Frey level, that knocks those guys out, right. Yes, teaching acoustics has done Dr. Swift a lot of damage in his career.

We can do a differential analysis. We can take an arbitrary unit-of-area slab of material and say that the heat into the slab minus the heat out is the heat that is left in the slab (eq. 43). Therefore, the net heat deposited or removed in this differential element of length, dx , causes the temperature of that element to go up or go down. This heat capacity, which I introduced in the first lecture -- this is the heat capacity per unit mass -- times the mass per unit of volume gives me the heat capacity per unit volume in this slug. That is the time-rate-of-change of the temperature (eq. 44).

We now apply the second most powerful technique in mathematics and we expand $\partial T / \partial x$, evaluate it at $(x + dx)$ in terms of Tdx at x and the derivative with respect to x of that quantity times that small infinitesimal distance dx (eq. 45).

If we substitute that result back (eq. 45) into this result (eq. 44), the leading term here cancels the leading term here is and we are left with this term. We end up with an equation of this form (eq. 46), which is not a wave equation. In fact, I am glad to say I have never written down the wave equation at this point. I still have not written down the wave equation. This is a diffusion equation.

If you reverse the direction of time and the direction of space, the equation changes sign, because eq. 46 has two spatial derivatives but only one temporal derivative. This combination of constants must have the units of meters squared per second. This has to have the units of temperature over time, and when you take the Laplacian, you are effectively dividing by two lengths, so the thermal diffusivity had better be length squared over time.

The difficulty is not velocity; velocity is length over time, or length squared over time squared. The diffusion constant has the dimensions meter squared per second. Therefore, it is usually given a number; that combination of constants is a thermal diffusivity, which I am calling χ (if that is what this Greek letter is called).

The form of this equation is important. This is a diffusion equation. It is the equation known as the Fourier diffusion equation; it is the equation for the diffusion of heat. It has got a first derivative in time and a second derivative in space. It is a very commonly occurring equation in science and engineering.

(Transparency #23)

It is isomorphic to the Navier-Stokes equation (eq. 47). The Navier-Stokes equation is the Euler equation that included the effects of shear viscosity. Here is the pressure gradient part, but here is the viscosity times the Laplacian of the velocity. So it is a diffusion equation for vorticity, or a diffusion equation for momentum. The next equation (48) is known as Fick's second law of diffusion.

If you have a concentration gradient, this is the change in concentration, and this is the diffusion constant, so it also has the units of meter squared per second. μ over ρ we call ν . That is also a diffusion constant. Its units are also meters squared per second.

Maxwell's equation in a conducting medium (eq. 48) is identical -- isomorphic, rather. The time derivative of the electric field is some combination of constants times the Laplacian of the electric field. This combination of constants is electrical conductivity and the magnetic permeability. Of course, their units are all over the place, because it is electromagnetism, but whatever this is, it had better have the dimensions of meters squared per second, because it is a diffusion constant.

I derived the diffusion equation for the simple case of thermal diffusion in one dimension (eq. 46), but it shows up everywhere. Any time you have something like Ohm's law, whether it be Ohm's law for electrical conduction or the law for shear forces and fluid flow or the law for thermal conduction, any diffusive process will have that form.

(Transparency #24)

The solution, for the time harmonic case, is an evanescent wave. As Phil Marston pointed out to me, the way I am using that term is correct but it is not unique; that is, in a waveguide you get an evanescent wave and its phase relations are not the same but it does have a characteristic exponential decay length.

Let us take thermal diffusion as a typical case. Let us say that we have a solid in contact with a fluid. I can modulate the temperature of this whole solid about some ambient temperature plus some modulation amplitude that is being driven at some angular frequency ω . We know that the diffusion equation is a linear equation. We know that the temperature in the fluid (here shown in the +y-direction to the right of that solid) must respond at that driving (modulation) frequency, because it is a linear system. And any driven linear system responds at the driving frequency.

Mark Hamilton will talk to you about nonlinear systems. They respond at the driving frequency and twice the driving frequency, at three times the driving frequency, and two times the frequency minus three times the driving frequency; all possible permutations. But in a linear system, if you drive it at ω , it responds only at ω .

So in the fluid, I will write that the fluid temperature, which depends on y , how far it is from this oscillating temperature solid surface, is T_0 , the ambient temperature of the fluid, plus some quantity T_1 , proportional to $j(\omega t - ky)$ -- so I am just assuming a wave-like solution.

If I substitute that form of the solution into the Fourier diffusion equation (eq. 46), the $\partial/\partial t$ gives me the $j\omega$, just as it did in the acoustic case, and the delta squared Laplacian (∇^2) gives me two products of a minus k . That gives me a minus $\chi k^2 T_1$ (eq. 49). If we want to plug that back into this solution, we can solve for jk (eq. 50) and we end up with the square root of $j\omega/\chi$. We know that j is a vector in the complex plane. As a vector pointing up and we know the square root rotates it by 45 degrees. So the square root of j has equal real and imaginary parts and, of course, the modulus is still unity.

So this jk is $(1 + j)$ divided by root 2 -- 45-degree angle, modulus 1 -- times the square root of the quantity in the parentheses (eq. 50). It is a complex quantity with equal real and imaginary parts. We can choose to define a real length, which is the real part of the inverse of that complex wave vector, which has the dimensions of length.

That is an important thing to do. In an acoustical system, we have a characteristic length for a system driven at a given frequency; it is the wavelength. For a diffusive system we have a characteristic length for a system driven at some frequency, and it is called the penetration depth. It is different from the wavelength in that it depends not on the frequency, but on the frequency to the one-half power.

If we substitute that result (eq. 51) back into my initial ansatz, we find the amplitude of the temperature oscillations decays exponentially with this characteristic length (in this case, the thermal penetration depth, δ_K) and it has equal real and imaginary parts, cosine y (the distance away from the surface over delta) and $j(\text{sine } y)$, and then there is, of course, the $e^{j\omega t}$.

STUDENT: I have a question. Can you show the previous slide?

(Transparency #23)

Here I think you usually write the equation in the Lagrangian fashion and not the Eulerian fashion?

DR. GARRETT: At the linear level it does not matter, because the $\mathbf{v} \cdot \nabla \mathbf{v}$ convective term has been thrown out. Thank you.

(Transparency #25)

That was a result that I derived for you algebraically (eq. 53). This is a plot of what you see. This is the case that I treated; that is, the wall temperature goes up by one unit and the fluid

at the wall -- the real part of that -- must be equal to the wall temperature. Across that solid-fluid interface, those fluids are in good thermal contact. The imaginary part initially increases as we go away from the interface and the real part decreases.

Acoustically, a slightly more interesting case is the one shown below (eq. 54), is that the fluid temperature far from the interface is oscillating by an amount, T_1 . The wall has infinite heat capacity. Therefore, the temperature at the interface is locked at the solid temperature and the fluid temperature is changing harmonically as it decays exponentially toward this isothermal surface. There is both an in-phase (real) and out-of-phase (quadrature) imaginary part.

Dr. Swift is going to show you all this in animation, which will make it a lot clearer, but I wanted to at least introduce that concept algebraically, so that you knew, one, where it came from and, two that it is ubiquitous.

(Transparency #26)

If we go back and look at not only the thermal boundary layer, but the Navier-Stokes equation, there is a viscous penetration depth (eq. 55). Again, it is that same combination of quantities, the square root of the diffusivity divided by the frequency. It is the kinematic viscosity. It has this characteristic dependence on one over root omega, which is, in some sense, the quantification of the old adage, "still waters run deep."

If you oscillate something slowly, that penetration depth goes to infinity. As you oscillate faster and faster, the inertia of the fluid keeps it from following the plate, but the stuff near the plate has viscous drag on it and it will move with the plate, so there an exponential decay, and that decay length gets longer and longer as you go to lower and lower frequencies.

In air, the length is roughly 2 mm divided by the square root of the frequency. It is about 100 μm at concert A, 440 Hz. The penetration depth in air for viscous shear is about 100 μm , 1/10 mm. The same is true for the thermal penetration depth, because the Prandtl number, the ratio of those two squared, is roughly one.

Here the penetration depth shows up again, but it is known by another name. This is called the skin depth in electricity and magnetism (eq. 56). In sea water the skin depth is about 30 m at 60 Hz. It is only about 2 cm at a MHz. That is why the entire state of Wisconsin is an antenna. If you want to communicate with submarines electromagnetically, you have got to go to very, very low frequencies.

Am I now under arrest, Logan (Hargrove)? Are those state secrets?

(Laughter)

Copper, of course, is a better conductor. At kilohertz it is only a couple of meters and at megahertz it is 66 μm . Exactly the same physics. The same thing is true with the mass diffusion (eq. 57), although they call it, again, by another name. They call it L . Instead of frequency,

chemical engineers usually like to talk in terms of diffusion times. But you get the same result: a random walk.

For propagation, as you wait longer, the distance increases linearly with time. For a diffusion process it increases with the square root of time. The common example is the drunk walking around the lamppost -- a random walk problem. Which reminds me of a joke that is rather important in experimental physics.

A gentleman had consumed too much alcohol and was outside a bar under a lamppost, feeling around, on his hands and knees, clearly inebriated. Someone else, much less inebriated, walked out of the bar and said, "Excuse me, sir, why are you on your hands and knees?" The drunk said, "Well, er, er, I dropped my keys and I can't get into my car." The other guy asks, "Well, can I help you find them?" (This, of course, is a bad idea, because this guy is drunk and giving him car keys is a mistake, but it is not part of the joke.)

The guy says, "I'd really appreciate that. I just can't find my keys." The other guy said, "Well, where did you lose them?" "I lost them in those bushes over there." "Well, then, why are you looking here under the lamppost?" "The light is better."

(Laughter)

You have got to remember that in the lab sometimes: It is fun to look, but it is not where you lost your keys.

For this diffusion time, these random walks, the distance increases not with time but with the square root of time. For argon diffusing through helium, it is about a centimeter times the square root of the time. If you start out with a concentration of argon and helium surrounding a point in space and He everywhere else during the first second, 1.2 cm away, 12 mm away, the concentration is $1/e$ ($=0.368$) of what it was when you started. Then, four seconds later, it is out to 2.4 cm. It goes like the square root; it does not double with time.

(Transparency #27)

That is characteristic of all of these diffusion processes, and it is important. The reason it is important is the answer to Dr. Hamilton's question: What is the speed of heat? We solved for ω and k . We know that the phase speed is ω over K . We have this result (eq. 49a) from the Fourier diffusion equation, assuming an harmonic excitation. We can solve for ω over K and we find this result (eq. 50).

It is dispersive, just like flexural waves in bars. Its velocity increases with increasing frequency. So it does not take a tremendous amount of mathematical genius to understand that if this is the speed of sound, a , and it is independent of frequency, and this is the phase speed of a thermal disturbance -- c^{thermal} , which is proportional to root ω , that at some point those two curves are going to cross.

Since the speed of heat is proportional to the root frequency, at very low frequencies the sound wave travels much faster than the associated temperature disturbance can. When you get a compression that is hot and you have got an expansion which is cool, the rate at which the heat from the hot part can travel to cancel the cold part is too slow.

By the time it has started to move, the phases have changed -- flipped back. That is why sound propagation is adiabatic. But there is a point at which the speed curves cross. At frequencies above this point, the sound wave would be isothermal, adiabatic below - isothermal above.

Since the sound speed does not depend on frequency, and this does depend on frequency, we can set them equal and solve for the critical frequency -- that intersection point. When you do that, you get this very simple result (eq. 51): γP over the heat capacity divided by the thermal conductivity. You plug in the numbers for air and it turns out to be half a gigahertz.

At low frequencies it is adiabatic and above half a gigahertz or so it is isothermal. That is almost right. The problem is that the wavelength at that point is under $1/10 \mu$ and the mean free path is $1/10 \mu\text{m}$. In an ideal gas you never get isothermal propagation, because the gas can no longer be treated as a continuum.

Our assumption that any characteristic length of the excitation sound wave contains so many atoms that we can treat it as a continuum just got destroyed, because we had to go to such a high frequency that in any given wavelength there is not a thermal equilibration process mediated by multiple collisions.

DR. BASS: Another issue here is you have ignored there, haven't you, viscosity?

DR. GARRETT: Absolutely.

DR. BASS: It makes the gas stiffer and when the gas is stiffer, the speed of sound is faster, and since the viscosity effect is larger than the thermal effect, you would never expect to see this.

DR. GARRETT: You will never see it, including viscosity makes it worse.

DR. BASS: Fluid viscosity overcomes the thermal conduction effect, so you cannot even see it in --

DR. GARRETT: Are you going to cover that in your molecular acoustics?

DR. BASS: I was not planning to. Do you want me to? I will mention it.

DR. GARRETT: It is interesting. I have just dealt with the adiabatic at normal frequencies, because it is pretty clear to do, but there are other interesting physical effects that we call visco-elastic transitions. The viscosity starts to behave like elasticity if you go to high-enough frequencies, and that is what you are talking about.

It is a wonderful area. It is an area that is very important in the physics of rubbers and relaxation processes for all kinds of things. All I wanted to do was quantify the situation at low

frequencies: Newton was wrong, Laplace was right. One can make this quantitative by using the simple models of the Fourier heat diffusion equation or the Navier-Stokes equation, as is appropriate.

But if you are down to 100 MHz, it is still not an issue. It is still adiabatic. When you get up there, things start to go bad. They are not going bad because you cannot get up there; it is going back because the assumptions that you made in doing these derivations to determine that frequency are no longer valid. You have stepped out of the regime of the applicability of those assumptions based on a phenomenological continua. It is an important thing to recognize.

(Transparency #28)

We can now discriminate between propagation processes and diffusion processes. In the absence of irreversible phenomena, sound is time-reversal invariant. Without dissipation, you cannot tell me how I made this lump of gas. I can say that I squeezed this gas together in some nice little Gaussian envelope there, and then let it go and these two pulses came out with the same width as the original lump.

If you say, yes, I believe that, I say, no, no, that was not the case. These arrows are wrong. These guys were actually coming together and time was going that way, and this happens to be the point at which they superimposed, and that is why it is a lump. You go, "Oh, yes, that's right."

Oh, no, fooled you. There is no way for you to win. No matter which way you say time is going, I can turn it around the other way and the physics will be the same. I like those arguments.

On the other hand, if we do the same thing, if we create a temperature distribution that is localized, in time it will get smaller and wider -- smaller and wider. Under no circumstances will a temperature distribution like this ever go to a temperature distribution like that. That is intrinsically irreversible.

DR. HARGROVE: You do not get shock waves in temperature in --

DR. GARRETT: That is Dr. Hamilton's area.

As we get to the shorter wavelengths, those gradients get higher and the short wavelengths attenuate more rapidly. Even though we have assumed perfect adiabatic propagation, there is some thermal conduction. If we put in a wave that has got "hair" -- like so -- after time the average amplitude will go down but the hair will go away. Proof over time that hair does go away.

(Transparency #29)

Given that we have these irreversible phenomena, like electrical resistance, there will be power dissipation. If we take a current, some sinusoidal current, and a voltage, and we want to know what the average power dissipation is, we take the product of those two and we average

over a period (eq. 53). We come up with a result (eq. 54) that the electrical power dissipated is one-half the amplitude of the current times the amplitude of the voltage, times the cosine of the phase angle between them. That cosine term is called the power factor.

The reason I write this up is that it is convenient to get the same result by using a complex representation. If we want to know the time average, we can take half the real part of voltage times the complex conjugate of the current (eq. 55).

This is just a way to get this result without having to do the integral. It makes life a little bit more pleasant. If you do that and calculate the viscous power dissipation, say, on some surface, that power is force times velocity (eq. 56). Work is force times distance; rate-of-change of work is force times rate-of-change of distance, or velocity.

I will define, following Dr. Swift's article on thermoacoustics, this power dissipation per unit area as \dot{e} and, therefore, the viscous per-unit area is \dot{e}_{μ} , which is the power dissipation due to viscosity per unit area, which is this (eq. 57). But we know the equivalent of Ohm's law, that force-per-unit area is p_{xy} , which is μ times dv in the x -direction, dy times V_x . That is the power dissipation.

Since we had the foresight to define the penetration depth as an exponential length as opposed to what the mechanical engineers do -- they define these distances as being 99 percent of the free stream velocity, for example, the Blasius layer, etc. -- if you make that mistake, the algebra is horrendous. If you use an exponential length rather than a 90-percent EE length or 99-percent ME length, this derivative is just v over δ .

If you plug this result back into this, you end up with a very simple expression for the power dissipation due to fluid friction per unit area is the viscosity times V_x squared over two times that viscous penetration depth (eq. 58), which is plausible. The power dissipation must be quadratic in the velocity, because dissipation is positive definite. It does not matter whether I am putting fluid over the plate going this way or putting fluid over the plate going this way; I am going to dissipate the same amount of energy. It is not going to depend on the sign.

Let us apply that to some acoustics that will be of use through the remainder of this week. This I think you have all seen at the Kinsler-Frey level. This is a Helmholtz resonator. It is a volume with a neck. Clint Eastwood was the mayor of Carmel. During his administration, the cops to stop you on the street and if you could not do this -- [the Eastwood theme whistle] -- you were under arrest.

(Laughter)

Anyway, you have a neck which has some fluid in it. That fluid has a mass, ρ naught A times L . That is just the volume of the neck times the density of the fluid in the neck, so that is a mass. We have solved the adiabatic gas law and we know that if those compressions happen, this wall can keep only a very narrow region -- thermal penetration depth, isothermal. In the

bulk of this volume the gas is adiabatic, so we will say, then, from the adiabatic gas law that $\delta p/p$ is $\gamma \delta V/V$, and the change in volume, of course, is just δx times this area (eq. 60).

Hooke's law of force, which is the area times the pressure, is minus kX . Therefore, equation 61 defines an equivalent stiffness, which I have solved here as just ρ times the speed of sound squared times a squared over V_0 (eq. 62). So if we know the equivalent mass, we know the equivalent stiffness, we know that the frequency squared of the simple harmonic oscillator is k/M (eq. 63). The resonant frequency of that Helmholtz resonator is proportional to the sound speed (eq. 64), as we would suspect (like any other resonator frequency).

It has some characteristic length, which is the length of the neck times the volume of the sphere, all of that over the area, so this has units of one over length squared. The square root of that is a length, so a/L dimensionally correct.

So that is the resonant frequency but, as my adviser used to say, after I had calculated various things to tell him the frequency, his next question was, "What's the Q?" I would go tromping back to my desk to answer that question.

(Transparency #31)

To answer that question, I would like to introduce a Helmholtz resonator that is currently being used at the National Bureau of Standards (sic) to measure the viscosity of gases and gas mixtures. It is known as a Greenspan viscometer, which is a double Helmholtz resonator. There are two volumes, a microphone and a loudspeaker, and a neck.

There is a reason you use a double Helmholtz resonator. If you are using gases and you do not want the thing open to the universe, you want to have a contained system. The slug of mass then oscillates between these two gas springs. When it moves to the right, it gets compressed here, which tends to push it out. It gets evacuated here, which tends to suck it in, so they both work in the same direction, and you get a resonance of that plug oscillating like so.

It is sealed, it has no radiation loss (I will talk about radiation loss in a minute). This geometry simplifies the transducer placement and it emphasizes viscosity rather than thermal losses -- I have to prove that to you. Resonance frequency, of course, is just that determined by the sum of the two stiffnesses (eq. 65) that are working together, and the equation (64) is modified only in that way to give equation 66.

You can see you revert to the open case (eq. 64). if you let V_2 go to infinity. You are back to my first expression for a single Helmholtz resonator, so it is not implausible that equation 66 is correct.

(Transparency #32)

What we want to do is calculate the quality factor to understand the dissipation. Why quality factor? Q is very, very useful. Number one, it is dimensionless. If I want to know what the typical sharpness of a resonance is, the size does not matter. Whether we are talking about

the crystal in my watch or talking about the balcony in a major rock-and-roll emporium that is oscillating like this, I can say what the Q is and you will know what I am talking about. It does not depend upon the size of the system.

It is also very convenient, because there are a whole bunch of different definitions of Q . If you want to calculate it in terms of 2π times the energy stored over the energy lost per cycle or energy lost per radian, that is Q (eq. 67a). If it is a lumped parameter system, we can define Q in terms of $\omega L/R$ for an electrical system, or ωM over the mechanical resistance for a simple harmonic oscillator (eq. 67b).

If we excite that harmonic oscillator and watch it decay in time, the $1/e$ exponential decay time is related to the Q , $\omega\tau$ over two (eq. 67c). If we are tuning a resonance, then the phase shift of the resonance, the rate-of-change of phase with frequency about the resonance frequency is related to the Q (eq. 67d). If we are measuring the half-power bandwidth, it is related to the Q (eq. 67e). If we are fitting this in the complex plane with a pole-zero plot, a resonance has two complex poles -- these are complex conjugates of each other -- the Q is related to the ratio of the magnitude squared and twice the real part (eq. 67e).

If we are looking at the elastic modulus of the materials, then Q is 1 over the dissipation, the loss tangent factor, or the real part over the imaginary part of the response function (eq. 67g). If we are looking at a damped system the Q is one-half of the critical damping ratio (eq. 67h).

So because there are so many definitions, it is very, very useful if we can calculate it in one case that we can apply it to free decay, to phase shift, to bandwidth, to dissipation, to critical damping. It is all there. So Q is something that we like to calculate right off the bat, because it is so useful.

If we look down the neck, it has got some area and there is this thin layer -- and I am assuming it is thin -- in which this dissipation is taking place. So the area is $2\pi r$, where r is the inner radius of the neck, times the length of the neck times the dissipation per unit area (eq. 68). Very straightforward. I can plug the result I had here that dissipation is $e\text{-dot}\mu$ is μ times v_x squared over 2δ (eq. 69).

(Transparency #33)

The energy stored, the total energy, is the sum of the kinetic and the potential energy. Because it is a harmonic system, because it is a resonator, that sum is equal to the maximum in the potential energy or the maximum in the kinetic energy. So I do not have to calculate the kinetic and the potential; all I have to do is calculate one of them.

Since the kinetic energy is $1/2 mv^2$ and I know the mass of the gas since it is the volume of the neck times the density of the fluid in the neck times the velocity (it is a gas slug that is all moving at the same velocity). I have got the kinetic energy (eq. 70) and, therefore, the Q is 2π

energy stored over the energy lost per cycle (eq. 67a). The energy is power dissipation times the period (the period is 1 over the frequency).

Plug the numbers in and you find you get this result (eq. 71), which simplifies, since μ over $\pi \rho f$ is δ^2 . You get a δ^2 on the bottom that cancels the δ^2 on the top and you come up with this extraordinarily simple dimensionless result: the Q of the Helmholtz resonator due to the viscous losses in the neck is just the radius of the neck divided by the penetration depth as long as that penetration depth is much smaller than the radius of the neck.

If it is not much smaller, you have to actually solve the Bessel equation. It turns out that Drs. Swift and Olsen have solved this for other reasons and there is a very simple expression which is good to 1 percent if the Q is greater than three.

So here (eq. 72) is the expression of Olsen and Swift for the Q without the restriction that δ/r is getting small. You can then correct the simple result (eq. 71) by these other two terms in (72) and it works fine down to Q 's as small as three. If you get into a bind, you have to solve the real problem, but in 90 percent of the cases the penetration depth for anything that has a reasonable Q will be small compared to the radius of the neck.

If we take the values for a single Helmholtz resonator, that is, a sphere about this long with a neck about this big and about that size, your standard sort of boiling flask, the Q turns out to be 58 for those numbers if I have plugged them in correctly.

If you go into Kinsler-Frey and you look at the expression for Q , you get 670 due to the radiation loss from that neck, so it is off by an order-of-magnitude. They do not write it in this nice dimensionless form; they have some really horrendous form. The point is that it is not even close. If you build a Helmholtz resonator and you open Kinsler and Frey and you do that calculation and you think you are going to get a Q near 1000, you are going to get a Q near 50.

The radiation loss is not an important factor in any Helmholtz resonator. It is usually the viscosity in the neck.

(Transparency #34)

We have not calculated the loss due to thermal conduction. The gas springs on both sides are adiabatic, so as the gas moves to the left it compresses this volume and it gets hotter here. There is an isothermal surface, there is irreversible thermal conduction. And similarly on the right. We have to worry about it, because the volumes have a hell of a lot more surface area, in general, than the neck. So why doesn't the thermal conductivity dominate?

We can go through and calculate the energy dissipation per unit area for thermal conduction and we come up with an expression (eq. 74), which I am not deriving for you, but that at least I will argue is plausible. If γ equals 1, it goes away. As we know from what I

did earlier, if a gas is isothermal, γ equals 1. If it is isothermal, there had better not be any thermal conduction losses, so it vanishes.

It is quadratic in the field variable. That is, when the gas is colder than the substrate or the volume, heat flows irreversibly in that direction and you have got dissipation. When you compress it, the gas is hotter than the substrate and it flows the other direction and you have got dissipation. So it has to quadratic, it has to result in a positive definite energy production, which it is.

Q is 2π times the energy stored over energy lost per cycle (eq. 67a). It is in my interest to express the energy stored as potential energy, since this is proportional to the pressure squared. The amplitude must cancel out in a linear system, it cannot be dependent on the amplitude. So I can express the energy density from the equation that I derived from conservation of energy (eq. 25).

I can integrate that for the volume, that is, I can multiply it by the volume, the pressure times the volume, and I end up with a Q (eq. 75) which is again plausible, in terms of what I told you before. R_0 is the radius of the sphere, $\delta\kappa$ is the thermal penetration depth, so you have that nice dimensionless ratio.

It is modified by a numeric here, because the ratio of the volume to the surface area is different for a sphere than it is for a cylinder. If γ approaches 1 the Q must go to infinity. If there is no temperature variation, there is no thermal conduction, there is no loss, the Q is infinite.

If we add up all the losses, $1/Q$ total would be $1/Q$ due to viscosity, $1/Q$ due to thermal conduction, $1/Q$ due to radiation. For that particular example you can see that the viscous losses dominate. In fact, in the Greenspan viscometer things are really chosen so that it dominates by even more than an order of magnitude. That is why it is called a viscometer. By measuring the Q of a double Helmholtz resonator you get the viscosity of the gas.

The question I pose to you is: if there is so much more area in the bulb volume and it is a surface effect, why is the Q due to thermal conduction so low? If you get to ask me questions, I get to ask you questions.

(Transparency #35)

I would now like to look at another Helmholtz resonator. It is a bubble. A bubble is a gaseous cavity in a liquid. That bubble obviously has a stiffness, because as we shrink the bubble down, the pressure in the cavity increases, which tends to push the liquid-gas interface back. The stiffness, again, is calculated using the adiabatic gas law, $\delta p / p$ equals $\gamma \delta V / V$. If we let "squiggle" be the virtual radial displacement of the interface, some small displacement, small compared to the radius, then δV is $4\pi R_0 \delta r$, the surface area of the bubble undisturbed, times that change in radial displacement over the volume, which is $4/3\pi R_0^3$.

Again, Hooke's law: force is pressure times the area and it is just minus K times whatever the equivalent coordinate is. We end up with an effective stiffness, which is $12 \pi \gamma P_0$ times R_0 (eq. 78). As you go deeper in the water, the ambient pressure goes up. Because you are going deeper and the stiffness gets larger.

The question is, where is the mass? If I am going to make a Helmholtz resonator, I have got the stiffness. Where is the mass? The mass is due to the fact that when the bubble grows larger, the fluid has to get out of the way. You have to accelerate the fluid away and then it accelerates as it comes back in. When you have acceleration of something, you have got mass.

It is very simple to calculate, although I will not do it. It turns out that Dr. Atchley is not going to do it, either. But the bubble frequency gives you sound wavelengths in the fluid that are very, very large. That bubble just pushes the liquid out. By Laplace's equation the velocity field has to attenuate, like $1/R$. If you integrate that velocity field, you will find that the kinetic energy stored in the fluid is equivalent to a mass of fluid that is three times larger than the mass of fluid which the bubble has displaced. That is not a universal result; that is a result for a fluid undergoing radial oscillations.

For a sphere-in-a-liquid-doing-this translation it turns out that the effective mass of the fluid that has to be accelerated to get away from the front or the back of this inclusion is one-half of the mass of the fluid that the bubble displaced. You have to do the calculation; that result is not universal. But it is simple in both cases and it turns out to be three times the mass of the liquid displaced by the bubble (eq. 79).

Therefore, the mass is given by the fluid that is being moved outside the bubble and the stiffness is the gas within the bubble and you get a resonance frequency (eq. 80).

DR. RASPET: Is that the density of water?

DR. GARRETT: Absolutely the density of water. That is why I wrote it ρ subscript H_2O . Why? Because you always forget. You go, oh, γp over ρ , that is a^2 . Bang, I have got it! Wrong. It is the pressure of the gas and the density, ρ , of the water. It is γ of the gas, it is ρ of the water. The natural tendency, Rich -- thank you -- is to do just what you did, is to say, oh, I know γp over ρ , it is the speed of sound in the gas squared.

(Transparency #36)

We can ask ourselves now what is the damping process? Again, Q is 2π times the energy stored over energy lost per cycle. We know the energy, we can express it as kinetic energy. We know the equivalent mass. We know the equivalent displacement, hence, the equivalent velocity. So that is $(1/2) m v^2$ (eq. 81). The radiated power is just like joule heating; it is one-half the radiation resistance times the volumetric velocity -- this is I_R , basically -- I have got it here, $RI^2/2$ (eq. 82).

This is the expression for the radiated power. It is dependent on the volume velocity squared. This is stored energy, dependent on the volume velocity squared. The result has got to be amplitude-independent, it is a linear system, Q cannot depend on the amplitude. You take the ratio and you get, again, this very simple and very elegant result (eq. 83), that the Q is just 1 over k times the radius, the k being the wavenumber: λ divided by 2π times the radius.

In this case we also have -- I do not have the bubble drawn any more -- the temperature of the gas changing inside due to adiabatic compression and expansion. The fluid has a much higher heat capacity than the gas, so it is isothermal at that fluid-gas interface. There are no viscous losses, because the motion is all radial, so there is no shear, and therefore we can calculate the Q thermal just as we did in the Helmholtz resonator case.

For a bubble of this size it turns out that they are rather comparable -- almost equal. In a Helmholtz resonator that you blow, the Q is determined primarily by the viscosity of the fluid passing through my thumbs and not by the radiation resistance, and the thermal conduction between the gas and the palms of my hands is negligible.

In the case of a bubble, the dominant Q 's are radiation, because it is a nice monopole radiator, and the thermal conduction losses.

(Transparency #37)

The point of all this is that irreversibility leads to a diffusion process, and even though there are all kinds of different diffusion processes, they are isomorphic, mathematically. They have the same form, they lead to the same result: thermal diffusion, or the Fourier heat equation; viscous diffusion, the Navier-Stokes equation; mass diffusion, Fick's second law; and electron diffusion, Ohm's law. They all have this characteristic length that varies inversely with the square root of the frequency, unlike sound, where the wavelength varies inversely with the frequency.

We now understand why there is adiabatic propagation: The hot parts of the wave are too far from the cold parts of the wave and there is not time to come to thermal equilibrium, thermal diffusion is not effective in the available time.

Adiabatic is only a limit. As you go to higher and higher frequencies, more and more of the energy is able to dissipate, but it is not sufficient to actually keep the temperature clamped.

Thermal conduction leads to dissipation, as viscosity leads to dissipation. This is the dominant effect, typically, in resonators. Boundary losses occur within this layer of thickness δ . The quality factor is a dimensionless measure of the loss. Energy ratios are a simple way to calculate that quality factor.

I hope you will carry that with you into the other courses that Dr. Swift will be teaching on thermoacoustics and Dr. Hamilton will be teaching on nonlinear physics.

DR. ATCHLEY: You have to be careful -- this has bit me a couple of times -- when you equate Q and losses directly. Just because a Q does go up it does not mean losses necessarily go down. It is dimensionless.

DR. GARRETT: Professor Bass is going to talk about classical attenuation, I assume, and the attenuation constant -- that is a poor choice, isn't it? -- is proportional to ω squared for classical processes. Since Q is $1/A \omega B$, or something like that, then the Q goes down with frequency -- sorry, it is ω over B . The Q goes down in frequency, because this goes up faster.

For boundary layer processes the dissipation goes up like root frequency. Therefore, as you go to higher frequencies, the Q goes up instead of going down. I think that is an important point that Professor Atchley brought out, that the Q is a dimensionless quantity, but it involves the frequency-dependence of the dissipation process which is important. It may be faster or slower than just the change in Q due to ωM over R_m as the simplest case.

Remember the lumped parameter version for a simple harmonic oscillator is ω , is going up with ω -- not a hard concept to accept. But the R_m , the dissipative part, may be increasing with frequency slower than this term is increasing and, therefore, the Q may increase with frequency instead of decreasing with frequency, even though the dissipation is increasing with frequency.

DR. MIGLIORI: You might as well mention that in solids, of course, you can go from the adiabatic low-frequency limit to the isothermal high-frequency limit in sound propagation, and right where you cross over the thing is as dissipative as crazy and it is not dissipating high or low.

DR. GARRETT: Dr. Migliori brings up an extraordinarily important point. In solids, the classic example is something known as the Zener effect. If you have a flexural bar -- like so -- above the neutral axis you have expansion. Below the neutral axis you have compression. So this is getting hot and this is getting cold.

But the characteristic length comparison is not the thermal penetration depth and the wavelength; it is the longitudinal thickness of the bar. So a bar can go isothermal at a very low frequency, because the distance between the hot parts and the cold parts is very small. It is not related to the wavelength. It is related to the thickness of the bar.

When you are in this region where you cross over, you get something that is known in Monterey, California as the Valdivia Sandpiper Model. For those of you who are not familiar with Jim Valdivia, he is a Naval officer. I was explaining all this to him and he said in class one day what Dr. Migliori is pointing out, "This dissipation that occurs at the crossover point is exactly like sandpipers."

I know that he was drug tested at regular intervals, since he was a Navy officer, and I did not smell anything on his breath, so I said, "Well, James, what do you mean 'it's like a sandpiper'?" He said, "Well, here in Monterey, I take my family out to the beach. Some days it's choppy. Some days it's calm. On the calm days the sandpipers stroll out with the tide and peck into the sand and the tide comes back and they stroll back in."

The sandpipers are relaxed, comfortable, they are not dissipating a lot of energy. On choppy days the surf is crashing at the same place. The sandpipers just stand there. They are not working very hard -- chop, chop, chop. He said, "But every once in a while I go out to the beach and the speed of the wave coming back and forth is exactly the same as the speed the sandpipers can run and you can see the sweat pouring from their feathers."

(Laughter)

So that became known as the Valdivia sandpiper model, because it is obviously correct. When you have a relaxation process, as Dr. Migliori is talking about, as you go from adiabatic to isothermal, the speed at which things are happening is matched to the speed at which things are dissipating and you get the maximum attenuation. This is a very common thing.

I think probably Professor Bass will be talking about this, right?

DR. BASS: Yes.

DR. GARRETT: If you plot the attenuation times the wavelength, it peaks at the relaxation point where you change from the sound speed being adiabatic to isothermal.

DR. BASS: It never goes in that direction, but you are almost right.

DR. GARRETT: Well, this is frequency times temperature, or something, so I can write it, right? It is -- bale me out, the Lendel-Ferry something-or-other transformation?

DR. BROWN: The William-Landau-Ferry?

DR. GARRETT: Yes, WLT transformation. This is a very important part for those of you who have academic careers: This is what is called the swift recovery.

(Laughter)

When you are wrong, you do not want to ever say that you are wrong. You always want to plot it in some other set of units and make you right, again.

(Laughter)

DR. MARSTON: Speaking of being wrong --

(Laughter)

-- we need to go back to your previous transparency on bubble damping.

(Transparency #36)

Even though you have radial motion in the water, you do have viscous losses.

DR. GARRETT: Excellent. How big an effect is it with regard to these other two effects?

DR. MARSTON: It is smaller than the effects at this particular size of bubble that you are talking about. But as you go to, I believe it is, the higher frequencies it can become more important. I think that was first discussed by Portis in the early 1950s and you can find it discussed in Devan's paper.

DR. GARRETT: Thank you. That is the classic paper on bubble attenuation mechanisms.

In my remaining five minutes I want to reward Albert Migliori for being so helpful in asking questions in the lecture at the proper time. He is going to be talking about resonant ultrasound spectroscopy, which is the use of resonances of physical bodies, of rigid solid bodies, to determine their material properties.

I am going to make just a few comments on that, because tomorrow, when you go to the Naval Postgraduate School for Professor Keolian's demonstration, you will be having a picnic beforehand at the school. I have set up a demonstration, that I am going to describe here, so that when you finish your lunch you can walk into the physics department on the first floor and my children and I have set up a nice little demonstration here on measuring the frequency of resonances in simple bodies, not the ones Drs. Migliori and Maynard use, to infer very accurate measurements of elastic constants.

I just point out that elasticity in a crystalline solid is not necessarily isotropic. There are preferential directions, or axes, and there are, in addition to compressive forces in a solid that are restored by the compressibility, shear forces. That is, a solid is an object that likes to maintain its shape, unlike a liquid, that takes the shape of its container.

Therefore, there are three components of strain for compressive forces and there are three components of shear -- shaking this way, this way, and this way. In principle, Hooke's law is now a tensor quantity -- this F (eq. 85) is not force, it is free energy, by the way. Hooke's law states that the stress (force per unit area) is linearly proportional to the strain, but since these are tensor quantities, the elastic modulus is a fourth-rank tensor.

DR. MIGLIORI: Actually, the tensor has 81 elements. When you introduce G_{12} , it goes U_{21} , you get 36. If you make it symmetric, you get 21.

DR. GARRETT: Right. What is commonly done, for those of us who do not like fourth-ranked tensors, is you treat this as being a vector. You call this $U_1, U_2, U_3, U_4, U_5, U_6$. Then you have the index running from 1 to 6, 1 to 6, and this is a 6×6 matrix, which gets you to where Dr. Migliori is.

DR. MIGLIORI: Yes, with not the right transformation properties.

DR. GARRETT: Right, but it is useful if you are doing piezoelectricity, because you end up throwing away a hell of a lot of the components anyway. But Dr. Migliori is entirely correct; you come to the same result. You get 36 elements. The symmetry about the diagonal reduces that to 21 independent elements.

If you choose your axes properly to align with the crystalline axes, you can reduce it to 18 independent elements. But that is more than K in $_KX_$ -- it is 18 times more complex for triclinic crystals or something.

For isotropic systems that are solids, this reduces further from 18 down to 2. For cubic symmetry it is 3. For isotropic materials it is 2. Then as it gets more complicated in its crystalline structure, it goes all the way up to 18.

But for an isotropic case, Young's modulus and shear modulus form a complete set. Two moduli are enough. We are back to phenomenology, again. Two moduli completely describe the system. If you choose, instead of Young's modulus and shear modulus, bulk modulus and Poisson's ratio, that is good, too.

But if you choose this and this, then the shear modulus can be expressed in terms of the Young's modulus. There are only two independent moduli. You can name all different kinds of moduli, bulk modulus, modulus for unilateral compressions, Poisson's ratio, shear modulus, Young's modulus. You can do that, but only two of them are independent. You, of course, are free to choose which two.

(Transparency #40)

If you have control over your sample, the moduli can be determined. If you have a sample like this, that is long and thin, so that the diameter of your rod is small compared to the wavelength, then the individual modes can be selectively excited and detected. We can detect and excite only the longitudinal mode, which satisfies the wave equation (86) -- see, I did actually give a three-hour lecture and I did actually get a wave equation on there -- where the sound speed here is longitudinal sound speed. It is the square root of Young's modulus divided by the density.

What you have are these drawings, courtesy of Keolian, who is not here, unfortunately. As the bar squishes in and bulges at the center and then stretches out and necks out at the center -- this is obviously an exaggeration, there are very few materials that would take that much strain -- but if you do that and you measure the resonant frequencies for these longitudinal modes (let us say the ends are free-free), then we can invert that, because the frequency is related to the wavelength, which is related to the length of the bar, and we can extract the elastic modulus (in this case, the Young's modulus).

(Transparency #41)

We can do a similar excitation of just the torsional modes; that is, we can take the bar and do this, then it is only the shear modulus that is the restoring force. If we measure the torsional resonances, we can get the shear modulus.

Similarly, the third mode that one can have is the one I drew on the board over there, the flexural mode. You can do the bar this way. It turns out to be a fourth-order equation -- shown

here (eq. 88). The constant here is, again, the same Young's modulus over the density, but the phase speed is now frequency-dependent, just as it was in the thermal case. At infinite frequency it does not go to infinite speed. There is something else, called a Rayleigh wave, that eventually takes over.

But at low frequencies it is proportional to the square root of frequency and, therefore, the modes are not integer multiples of the fundamental, whereas in the torsional case (eq. 87) and the longitudinal case (eq. 86) it is a second-order wave equation. For free-free boundary conditions with no mass loading, no problems; they are harmonic.

(Transparency #42)

I have summarized the theory here to illustrate that if you have a longitudinal wave then the elastic modulus, Young's modulus, in this case, depends on density, which you can measure to high precision, and length, which you can also measure to high precision, and the frequency divided the mode number squared.

If you measure the frequency and you know something about the sample, you can determine the Young's modulus. The same thing is true for the shear modulus. If you know the density of the material, the length of the sample, and the various frequencies, you can determine shear modulus if you have a system where those modes are separable.

Similarly for flexural modes, except that the mode numbers are no longer simple integers until you get to fairly high mode numbers, because it is governed by a fourth-order equation.

(Transparency #43)

What you will see tomorrow at the Postgraduate School is a little rig that we have been using for a long time that supports a rod. We wind up a coil of copper wire, usually around the butt of a screwdriver -- it has those little holes in it and then you can weave a piece of wire there to hold the little coil together -- and then you squinch it down on the rod, as shown here, and glue it into place.

If you place that coil within a magnetic field, as shown, B is this way, the current in this leg is coming out of the board, the current in that leg is going into the board, and that gives you a net torque. If you change the sign of the current, it torques it the other way. That selectively excites the torsional mode.

Similarly, at the other end, there is the same setup, and because you have basically this galvanometer coil between these two magnets, as you tilt it, the flux changes. This way there is zero flux, this way there is positive flux, this way there is negative flux. Faraday's law: You are generating voltages due to this flux through the coil changing and you selectively excite and detect only the torsional mode.

Here are some data, which are actually fairly typical data. This is not a great sample; it is a plastic. With metals the Q is 10 to 100 times higher. But you can see that the signal to noise

here is in excess of 70 dB. First mode, second mode, third mode, fourth mode, fifth mode: They are harmonically related; that is, amplitude is plotted versus frequency, the spacing between the resonances are equal.

(Transparency #44)

If we take the rod, rotate it by 90 degrees and displace it vertically as shown, we can excite and detect preferentially the flexural mode, because the magnetic field on this part is stronger -- it is between the pole pieces -- and the magnetic field on this portion of the wire is weaker, because it is in the fringe field.

Therefore, there is a net force pushing this up that is greater than the force on this pushing it down, and it flexes the bar. What you see here, again (this is a plastic sample, not a metal sample), is the first three resonances, and you can see the spacing between this and this is roughly three divisions, whereas the spacing between this and this is four-and-a-half divisions. So they are not harmonically related.

But again, now you see, going from 2 dB to -78 dB, there is roughly 80 dB of signal-to-noise. The measurement of this frequency is not difficult, that is what I am trying to point out. You will see it set up in a little gizmo in the lab at NPS. Similarly, you can do the longitudinal mode, but because there are only two independent moduli, Young's modulus and shear modulus, then this provides a redundant determination of the Young's modulus.

These are longitudinal modes, but you can see here the first mode is quite strongly excited but the second the third modes are somewhat weaker, because now you have the magnetic field and it is only the current in this short section of wire at the end that is creating the force and changing the flux.

With that, I hope that I have taken some of the issues that you are going to be seeing in this afternoon's lecture by Professor Bass on molecular acoustics and related it to the phenomenological models; we have looked at irreversible processes, heat transport, energy transport, Q's of resonators, and now some introduction to the relationship between the elastic moduli of solids and their determination by resonant excitation of various modes.

MOLECULAR ACOUSTICS

HENRY BASS

UNIVERSITY OF MISSISSIPPI

DR. BASS: This afternoon I am going to take a quite different approach than was taken by Dr. Garrett in his discussion this morning. Specifically, he began to talk about acoustics in terms of a phenomena, in terms of the macroscopic behavior of the system, in terms of measurable quantities.

His approach is reasonable, because Dr. Garrett is an experimentalist so he relies on experiment to guide him in his selection of a description of the system. In fact, if we look back in the history of physics, that is the way physics has evolved. It has been through experiment that we have observed different phenomena. These observations have then led to theoretical descriptions and models.

Physicists have traditionally tackled problems this way; they have gone to the lab, they have done experiments, they have come out with results that they do not quite understand, and then they try to formulate models, microscopic models sometimes, that help them understand those results.

This is also true of sound propagation in gases. This afternoon I am going to spend most of my time talking about microscopic acoustics, or molecular acoustics, talking about the properties of sound in gases and how the microscopic properties of gases affect that sound while it is propagating.

This field started out as a result of people making careful observations. One of the careful observations Dr. Garrett explained to you had to do with the speed of sound. When Newton first formulated his speed of sound, he was wrong. He was wrong because he left out a factor of gamma. He assumed the sound speed was isothermal rather than adiabatic.

In order to come up with a better description, he needed to know something, or Laplace needed to know something, about the nature of the media through which the sound wave was propagating. I will get back to that in a moment.

Also important to our understanding of sound were early measurements of sound attenuation. As sound propagates outward from a source, there is, of course, spreading that causes the amplitude to fall off with distance as $1/R$.

But early experimenters, even back near the turn of the 1800s, recognized that the sound amplitude decreased more rapidly than just $1/R$. Other factors are involved. There were attempts to explain these, but there was no good experimental data. It was just reported that sound was being attenuated more rapidly than expected.

As an example, in a fog in England, where these experiments were taking place at the time, sound propagates much differently than on a clear day. There was obviously something taking place other than just $1/R$ (spreading losses).

Much of that has to do with vertical temperature gradients in the atmosphere and had nothing to do with the theories that were formulated at the time, but that is beside the point.

Stokes and Navier -- Navier-Stokes -- entered into the picture and they postulated that sound as it propagates was also attenuated, absorbed, due to viscous effects, due to viscosity. It was not a very large step beyond that, when Fourier's law was introduced, to include losses due to thermal conduction.

Absorption due to viscosity and thermal conduction is typically referred to as classical absorption.

(Transparency)

I will start out talking about classical losses where Dr. Garrett left you, at the phenomenological level. I will describe absorption of sound in gases due to viscosity and thermal conduction.

Then we will go on and examine absorption in even greater detail and find that viscosity and thermal conduction are not sufficient to describe the attenuation and dispersion that is measured in most gases. This really was not well documented until after the turn of the century. One of the people who was important in this process was Keneser in Germany (and Knudsen in this country), who recognized that the attenuation that was observed in molecular gases was greater than predicted by the Navier-Stokes relation.

This was about the time that quantum mechanics was getting a full head of steam. In fact, the first careful measurements made of sound velocity dispersion in gases as well as absorption in gases were used to demonstrate that gases, in fact, had internal modes of energy which were generally accepted, and that these internal modes of energy were quantized.

Careful measurements of absorption of sound in air, showed that the quantum mechanical description of vibrational energy in gases required a zero-point energy, something that was only hypothesized at the time. So the absorption and dispersion measurements in gases contributed to the development of microscopic models of gases.

After 50-60-70 years we accept these microscopic models as being the base from which we start. But keep in mind that it was the disagreement between the predictions of prevailing theory and measurements that gave rise to the discovery of some of these very important concepts.

I start talking about classical absorption and then go on to relaxation absorption. Relaxation absorption involves internal degrees of freedom, internal structure. This internal structure and these internal degrees of freedom manifest themselves in added absorption of sound and dispersion.

Finally, I am going to talk about diffusion and absorption at low pressure. I am going to add dispersion since the subject came up this morning; what little I can contribute. And then the low-pressure regime. Some interesting things begin to happen when you try to propagate sound through a gas in the very low-pressure or the very high-frequency regime because you reach the point where the acoustic wavelength becomes comparable to the mean free path.

When those two quantities become comparable, many of the assumptions that we quite often invoke when describing acoustic phenomena break down.. This does not mean that acoustics does not work any more; it simply means we have to go back and reexamine some of the assumptions. We will do that in about 20 minutes.

I might warn you now that when I get to that point, I do not have many answers. I have questions. One of the reasons we do not have any answers is because there has not been a lot of work done in this area in the past 30 years. Since Moe Greenspan died, this area has gone unnoticed. It is a shame; maybe some of you will fill in this gap. Greenspan's experiments ended about the time that solid-state technology was introduced into the physics laboratory.

We now have means of dragging signal out of noise that are far superior to anything that was available to Greenspan. We have a whole new generation of measurement abilities that has yet to be brought to bear on this problem, and there are some interesting questions that he left unresolved upon his retirement.

Let me point out, as did Dr. Garrett, that I have given this talk a time or two. I have rehearsed it, and I know that I can give the talk in three hours -- or, actually, three hours, including two 10-minute breaks. Dr. Hamilton is laughing; he has heard this.

We all prepare ourselves, because we are nervous if we are not prepared. But if because of questions or because you want to stop me and ask, "How about this," we can throw t away what I have prepared. I have written out what I am going to say and it was sent to you ahead of time. So I do not have to cover the prepared material. If we get off into something else that is more interesting to you, that is fine.

I am up here to tell you about some things that I think are interesting. If you want to go off in different directions, we can do that.

DR. GARRETT: Dr. Bass, it is reputed that there is a model board at Ole Miss. Could you make some comments about that, instead?

DR. BASS: There is a modeling board at Ole Miss. I am not sure this is a politically correct subject to discuss. Dr. Garrett has a history of bringing up politically incorrect subjects.

(Laughter)

At Ole Miss we actually have what we call sensitivity training. If someone is found to more than once bring up a subject that is not politically correct, we send him to sensitivity training. That can last anywhere from four hours to 40 hours each time around. Maybe Dr.

Garrett is a candidate. We have an office called "Politically Correct Police" as well, but that is a different story.

We have a modeling board. Ole Miss has a tradition of having some of the best-looking young ladies in the United States attending our university, something we are very proud of. Some of these young ladies are also extraordinarily bright. Some them are not. That is just the way it goes.

(Laughter)

We have produced as many Miss Americas as any university in the United States, another record that we are proud of. We have certainly produced more Miss Americas than we have Nobel Laureates.

(Laughter)

We have only one Nobel Laureate to our credit. We have several Miss Americas to our credit.

The modeling board is the launching point for some of these young ladies to begin their careers. They actually have sessions where they learn to model clothes and they learn to walk properly and they learn to hold their heads properly and they learn to prepare their hair properly.

This is a very important part of their education. Quite often this training lands them a very wealthy husband. But sometimes they get to be Miss America. So yes, we still have it.

MS. CAUTHEN: How does that tie in with molecular acoustics?

(Laughter)

DR. GARRETT: Everything on molecular acoustics is in the handout. He said we were free to talk about other subjects. I had an interest in another subject.

(Laughter)

DR. BASS: I opened myself up and I agreed to address any topic that was brought up.

DR. GARRETT: Thank you.

DR. BASS: You are welcome.

(Transparency)

Let me take just a few minutes here to reiterate some of the things that Dr. Garrett said. Let's start with an acoustic wave equation. An acoustic wave equation is derived from conservation of mass and conservation of momentum. As Dr. Garrett pointed out, if we treat conservation of momentum as a vector quantity, we need no additional equations.

Or if you do not want to treat momentum as a vector, you must include conservation of energy -- that is a separate equation. Next assume some acoustic form for a solution which allows the pressure to vary sinusoidally in both space and time. Solving these equations together gives an equation for the speed of sound.

The equation for the speed of sound that I have written here is the Laplace equation. It has a γ in it and p naught over ρ naught which, according to the ideal gas law is RT/M . So we get that the speed of sound is equal to the square root of $\gamma RT/M$, which is the same as Dr. Garrett's result.

In this case we had no absorption mechanism, no dissipation, so the sound velocity is pure real.

(Transparency)

Now, include some dissipation mechanisms, some attenuation. First, include viscosity. Refer to the top equation on the transparency. The left hand side is acceleration. The right-hand side force. I will modify the momentum equation to include not only the partial of p with respect to x , the gradient of the pressure and a new quantity. That new term is noted by viscosity, η -- this should be η prime.

The $4/3 \eta$ part is the term that we are typically interested in. The second term, η prime, is referred to as the bulk viscosity. Bulk viscosity was introduced by Navier because he was not really sure that the other term, namely, the shear viscosity, would tell the whole story.

In fact, it does not. There are other processes that one can lump into bulk viscosity. But traditionally we have chosen to put those other terms somewhere else and to treat them separately. Following this traditional approach, we set the bulk viscosity, η prime, equal to zero.

Bhatia, from India wrote a book in which he included all relaxation processes, microscopic processes, in this η prime term. But I will not do that. I will set η prime equal to zero.

That leaves me with only a shear viscosity. If you stop and think about that for a moment, you should get confused -- or at least I get confused. Think about a plane wave propagating in the x direction, then the gas squeezes like this.

If the gas squeezes like this as the wave propagates through, why should a shear viscosity term enter? Where is the shear? You figure it out. In any case, if you assume that shear viscosity does enter then, as I have pointed out in the handout in a little more detail, we must assume a slightly different form for the solution. Again, it is sinusoidal, but now it has a different term in the exponential because the solution is complex.

In this case, then, I get not only a propagating part of the solution, which I saw for the non-loss case, but also an attenuation, a term that causes the amplitude to decrease exponentially with distance.

The important thing that I want you to note here is that this attenuation, this absorption coefficient, that appears in the exponential depends upon the frequency squared.

Dr. Garrett noted this behavior. He pointed out that we see this form cropping up repeatedly in other terms that are introduced into the wave equation. He had a fancy name for this.

If I put in some numbers absorption in oxygen is about 10^{-8} (this is gaseous oxygen) times the frequency, and for water it is about 10^{-12} times the frequency. The absorption due to viscosity is about four orders of magnitude greater in a gas than it is in a liquid, and for most other absorption mechanisms that continues to be the case. The term in the exponential decay is about four orders of magnitude greater in the gas than in a liquid.

That is why one can propagate sound from Heard Island to the coast of the United States through the ocean but if you try to shout at your kids two blocks away, the sound does not get there.

I write here that the speed of sound is almost unchanged by the presence of viscosity; in other words, the sound speed in the presence of viscosity is almost equal to the sound speed in the absence of viscosity. The "almost" is what we were talking earlier this morning. I will come back and address this a little later. But let me skip over that for just a moment --

(Transparency)

-- and remind you of another source of loss, heat conduction. Dr. Garrett also discussed this term. We can include heat conduction using Fourier's law, which says the time rate of change of heat flow depends upon the temperature gradient.

I want you to note that here, again, is the dependence on omega squared. I want to come back and dwell on that issue a little later.

Let me now take a little diversion to address dispersion. This equation says that the velocity is the same in the presence of thermal conduction as it is without thermal conduction. That, as Dr. Garrett pointed out, is not quite true.

(Transparency)

I do not agree with Dr. Garrett that the speed of sound decreases at very high frequencies. I did a little doodling a while ago and came up with the following for a monatomic gas. If I take the Lagrangian speed of sound and divide it by the measured speed of sound, the true speed of sound should be in this ratio.

There is an expression that relates thermal conduction and viscosity. Thermal conduction and viscosity are both a result of the collisional transport of energy and momentum from one gas molecule to the next. When a collision takes place, there can be a transport of momentum. But as you know from sophomore physics, if you transport momentum, you are also going to transport some energy. You have to conserve both at the same time when you have a collision; you cannot do one without the other.

That might lead you to believe that there is a relationship between thermal conduction and viscosity and, in fact, there is -- there are a number of them, depending on who you listen to. But no matter which one I choose the relationship between the coefficient of thermal conduction and the coefficient of viscosity looks something like the one on this transparency.

Using this relation between viscosity and thermal conductivity I get an expression for the speed of sound divided into the Lagrangian speed of sound. My back-of-the-envelope stuff says that the measured speed of sound is always going to be greater than the Lagrangian speed of sound, which is always going to be greater than the Newtonian speed of sound.

Note that this term is going to be largest, so the difference will be greatest, when the frequency is high. The higher the frequency, the greater is C and the larger it is with respect to the Lagrangian or Newtonian speed of sound. I also have some experimental results that support this conclusion.

(Transparency)

These are experimental measurements made by Mo Greenspan. The measurements I refer are at the top. The theory is the solid line, the measurements are the little points. They do not agree that well; the reason comes later. But note that the ratio shown decreases at very high frequency. C_0 is the low-frequency speed of sound. C is the measured speed of sound. Note that the measured speed of sound gets larger at higher and higher frequencies. The gas gets stiffer, in other words.

This sort of makes sense, doesn't it? Let us think now in terms of the gas as a microscopic medium. Energy is transported by the flight of molecules from one point to the next, where they collide with something and transport information.

At very low pressure or very high frequencies the molecules do not collide as frequently between acoustic cycles. If you think about it for a moment, you would probably agree that the important quantity is the ratio between the mean free path and the wavelength. I can change that ratio either by increasing the mean free path, which means to lower the pressure, or by decreasing the wavelength, which means increasing the frequency. Increasing the frequency is equivalent to decreasing the pressure.

At very low pressures or very high frequencies the molecules do not collide as often. In fact, you can even imagine the case, and we will get to that a little later, where molecules start out at an acoustic source and get to a microphone without any collisions -- well, then, the velocity of propagation would be the Boltzmann velocity in the gas.

The Boltzmann velocity, depending on who I asked at lunch (I did not get a consensus), is anywhere from three to seven times the speed of sound. No? How much is it, Dr. Garrett?

DR. GARRETT: It is only about 20 percent higher. I will work it out and tell you the answer.

DR. BASS: Okay. It is higher.

DR. GARRETT: Yes, but it is only by a few percent.

DR. BASS: But it is higher. So in the extreme case of very high frequency, molecules will travel with the mean speed for the molecules, which is greater than the low frequency speed of sound. We would, therefore expect this curve to turn down; it does. The question is whether it turns down by 20 percent or by a larger number; Dr. Garrett is going to tell us what the number should be. Does that make sense to you?

DR. MIGLIORI: You mean about the stiffness?

DR. BASS: Yes.

DR. MIGLIORI: Yes.

DR. BASS: In the gas there are two effects. One is thermal conduction, which would have the effect --

(Transparency)

-- of decreasing the speed of sound. This $4/3$ comes from the viscosity term, which has the effect of making the gas stiffer, increasing the speed of sound, but since this term is larger than that one, it wins and the stiffness increase overcomes the thermal conduction decrease.

DR. MIGLIORI: That is primarily a geometric effect for a plane longitudinal wave at this point, the fact that the cancellation goes --

DR. BASS: I think that this is the result any time you have that form for the expression.

DR. MIGLIORI: Yes, that is basically the Weidemann-Franz law for gases, right. So this is a result, really, of a 1-D solution to the wave equation.

DR. BASS: This is not dependent on the dimensionality of the solution to the wave equation. It depends only upon the solution to the collision problem.

DR. MIGLIORI: What if you did it for a surface acoustic wave?

DR. BASS: I have no idea.

DR. MIGLIORI: Yes, see, then you may have a different -- I do not know.

DR. BASS: I have no idea. I think, though, that one of the important issues is that this ratio does not hold for other than an ideal gas.

DR. MIGLIORI: That law, when you do it in a fully quantum mechanical system, including all of the details, energy-dependent scattering, nonclassical dispersion curves as, for example, in an electronic system, that law actually winds up holding extremely well, unless the relaxation time approximation is killed.

That is the one that says when you have a collision, you basically reset all of the information about the velocity equation.

DR. BASS: Let us now introduce a part of this problem that Dr. Garrett did not describe in great detail.

(Transparency)

Let me start off by describing in a hand-waving way the phenomenon that is going to be of interest to us. That has to do with the transfer of energy during collisions at the microscopic level. I pointed out earlier that the way a sound wave propagates from point A to point B is through collisions. You disturb a molecule here, it collides with its neighbor and transfers information to its collision partner. Through collisions, information is transferred down the line. That is how the sound wave propagates.

As the wave propagates from point A to point B, there is a local disturbance; an increase in pressure and temperature. The increase in temperature causes the gas to increase the energy in each of its available modes.

Dr. Garret described to you the equipartition of energy. Equipartition of energy, the remnants of communism, is at work not only between the different translational modes, X, Y, and Z, but also at work between the various other modes available for energy to the molecule. For example, a nitrogen molecule can not only translate, it can also rotate.

Think of the molecule as a dumbbell. If it rotates, that rotational motion has energy. Not only can nitrogen rotate, it can also vibrate; it has a chemical bond between two nitrogen atoms and they can go like this -- [hand gestures]. So it has another possible place to put energy.

The rule that everyone shares equally says that if you locally increase the temperature through collisional processes, the first thing that is going to increase is the translational energy, because that is mode which leads to the collision. But very soon that energy is going to be shared with rotation. Whenever you have a collision most likely going to be slightly off center which will cause the rotational energy of the molecule to change as well.

For atoms, the situation is different. Atoms are little round hard balls. You can think of them as little tiny billiard balls. Everything there is really cool, because all you can do is translate vertically and horizontally, so you have only translational degrees. In this case everything that Dr. Garrett told you is true.

But most gases are not made up of monatomic gases.

We have these other modes which we must consider; primarily, rotational and vibrational energy modes.

At standard temperature and pressure there are about 10^{10} collisions per second. Transferring sound from one point to another involves a large number of collisions.

The transfer of translational energy is like billiard balls running together. If one billiard ball is stationary and all the billiard balls around it are going 50 m/sec in random directions, it is not going to take very many collisions for the stationary billiard ball to come up to speed. Within a few collisions the billiard ball that starts out at rest should come up to speed with the rest of them.

In fact, one can go through a calculation and find out that in order to come up to $1/e^{\text{th}}$ of the speed of the surrounding molecules, it takes about 1.25 collisions (this is a fairly difficult calculation and I will not try to justify it). My point here is that in order to reestablish equilibrium for translational motion, once lost, little time is required.

If it takes 1.25 collisions and there are 10^{10} collisions per second, it is only going to take a couple of nanoseconds at atmospheric pressure in order to reestablish translational equilibrium in the gas. For all practical purposes, translational motion can be thought of as reacting instantaneously to an imposed sound wave.

Rotational modes are almost as fast. The thing about rotational modes that one has to be just a little careful of is that rotational energy for molecules is quantized.

For most molecules that is no big thing. For a molecule like nitrogen the rotational energy-level spacing is very small compared to kT ; almost all collisions have energy far in excess of that required to cause a Δj equals 1 transition. So there is not going to be a problem causing the rotational state to change.

Experimentally, we observe that to be the case. In nitrogen, it takes on the order of five collisions to reestablish rotational equilibrium once it is lost. But there are some exceptions: hydrogen, for example. In the hydrogen molecule you have these two little bitty atoms at the ends of a real long stick. The moment of inertia is much different in this case.

The rotational energy-level spacings in hydrogen for energy levels occupied near room temperature lie 50-100° K apart in energy units. You get Kelvin energy units if you multiply $h(\text{Planck's constant})$ times ν , where ν is the frequency. The Boltzmann constant times temperature gives an effective energy of translation. This terminology allows me to talk about energies and temperatures interchangeably.

The point here is that for hydrogen many collisions do not carry in enough translational energy to cause a Δj equals 1 transition. There is not enough energy available to cause the rotational quantum number to change.

Those collisions that do not have enough energy to cause the rotational quantum number to change are not going to be effective in reestablishing rotational equilibrium. In fact, we find that for hydrogen the number of collisions required to reestablish equilibrium, once it is lost, is on the order of 350 near room temperature.

If you are interested in studying hydrogen or other gases at very low temperatures, this can be important. In most cases we are going to be interested in other molecules or operate at lower frequencies, so we often make the assumption that rotational relaxation also takes place almost instantaneously.

MR. HALLEJ: When they have the collisions in hydrogen, energy is going to get transported anyway, but it is just not going to make its way into the rotational. So they will keep

their energy in translational, right? So you are saying that hydrogen will have a different distribution, or it is not going to be equally partitioned between translation -- it will have more in translational, so the velocity of a hydrogen molecule will be expected to look different from nitrogen?

DR. BASS: The translational energy distribution will be Boltzmann for both of them. The rotational distribution is not going to be the equilibrium distribution for hydrogen. If you are in a part of the sound wave where the temperature is going down, then there might be more energy in the rotational modes than equilibrium conditions would dictate. If the temperature is going up, there might be less energy in the rotational modes than equilibrium conditions would dictate.

The point is that the effective temperature of translation is different from the effective temperature for rotation.

DR. HALLEJ: What is important for the sound wave is translation, isn't that true?

DR. BASS: The answer to that question is no, but we will see it in one more slide. Do you want to try to clarify that, Dr. Migliori?

DR. MIGLIORI: It is only in the dynamic processes that the effective temperature of the rotational modes is different from the effective temperature of the translational modes. So there are frequency ranges in which complete equilibrium is the correct approximation and others in which complete equilibrium of the translational processes is, and so on. So you get lots of little - but you are going to tell them all this.

MR. POESE: When you talk about a monatomic gas you talk about a billiard ball sitting still and a bunch of other billiard balls hitting it. It takes 1.25 collisions for the single billiard ball to get back up to equilibrium. What about if I make one of the balls move a lot faster than the rest of the balls? I mean, it has to be the same thing, right?

DR. BASS: Yes.

DR. POESE: Where did that energy go, then? I am having a hard time understanding how --

DR. BASS: It would go into increasing the speed of all of those other billiard balls in the system.

DR. POESE: I see, so the same thing with rotation? If I can put a lot of energy in the system and that is enough to get something to start to rotate out of rotational balance or out of rotational equilibrium, then everything is still going to rotate faster, is that it?

DR. BASS: That is right. Let us suppose that you put energy into translation, you increase the local temperature, which means that the sound wave is at a maximum in temperature. Rotational energy does not quite keep up. It will given enough time.

When the energy of the translation goes down, the rotational energy is too high, so it will start dumping this energy, but it will be dumping it a little bit too late. What is going to happen?

You are going to be taking this energy out of translation and putting it back into translation just a little bit out of phase.

MR. PRATHER: So the rotation....

DR. BASS: And what is going to happen to the sound wave?

MR. PRATHER: It attenuates.

DR. BASS: It is going to attenuate.

MR. PRATHER: Right.

(Transparency)

Let us see how that works in a computational way. I do not want to deal a great deal with equations, but I do want to show you how this failure to maintain equilibrium enters into the wave equation. First let me point out that in addition to rotation there is vibration. The only difference here is one of microscopic description -- vibrational modes look different, instead of doing this, they do this -- [hand gestures]. The number of collisions necessary to reach vibrational equilibrium is many orders of magnitude greater than required for translation or rotation.

When nitrogen in vibration --going like this -- collides with another nitrogen molecule, an argon atom, or whatever, to reestablish a new vibration equilibrium will take on the order of 10^{11} to 10^{12} collisions.

The reason for that is as follows. The vibrational modes, the quantum mechanical vibrational modes in nitrogen are spaced at great distances. In fact, if you assign them a temperature, that temperature will be about 3300° . If you examine the distribution of velocities of colliding molecules, only those few that have energies corresponding to a temperature of 3300° can cause such a transition.

In the Boltzmann distribution most of the molecules have a temperature near 293° at room temperature. Only those molecules way out on the tail of the distribution-- remember, that tail goes something like $e^{-\text{exponential}}$ minus velocity squared -- are going to enough energy to cause a vibrational transition. Very few collisions are going to carry in sufficient energy to allow a vibrational transition to take occur.

The number of collisions necessary to cause a vibrational transition in nitrogen is about 10^{11} , which means that it takes about 10 seconds, roughly, once you disturb nitrogen, to come back into vibrational equilibrium. There are other ways to achieve an equilibrium distribution.

For example, in collisions with water vapor molecules, rather than transferring energy directly from vibration to translation part of the energy can go into rotation. Let us say you want to excite a 2000° energy gap. If you have three collisions and each of them increases the rotational temperature up by 700° and then that rotational temperature transfers to vibration, you have a much higher probability of getting to 2000° than by doing so all in one step.

There is another way of beating the system. That is through a vibration-to-vibration process. Consider water vapor molecules colliding with oxygen. Water vapor has lots of vibrational degrees of freedom. One degree of freedom is the asymmetric stretch. The vibrational mode of oxygen is almost in resonance with this asymmetric stretch of H_2O . So transferring energy back and forth between those two is very easy, because there is almost no energy deficit in the collision. Most collisions will have sufficient energy to cause a transfer of energy from water vapor to oxygen.

Once the energy gets into water vibration, it cascades down through rotational states, so it relaxes very rapidly. If you put a little water in with the oxygen, the oxygen transfers energy by vibration-to-vibration transfer to H_2O and H_2O communicates with translation. Water vapor catalyzes the reaction.

Vibration to vibration transfer tends to take place very rapidly, perhaps with a collision number of only a few hundred. Vibration to translation transfer might take 10^{11} collisions. There is a wide range in the number of collisions required to cause a transfer of a quantum of vibrational energy. Rotational collision numbers vary between about 4 and 350. In the case of vibration it varies from about 10 to 20 to 10^{11} . No one has successfully measured the vibrational relaxation time of molecular hydrogen, H_2 . It could be, and probably is, measured in years.

MR. PRATHER: If the relaxation time gets long enough you said earlier that energy comes back out of phase resulting in attenuation. But if the relaxation time is very long, doesn't it become random, again?

DR. BASS: That is a good point. If you delay putting energy back in by a small amount, maybe 90 degrees or so, you get attenuation. What happens if you delay it even longer? Remember this curve that Dr. Garrett put up on the board earlier (plot of α lambda versus frequency).

Basically, that particular mode is frozen. It can no longer communicate with changes in translational energy, so it is like that mode does not exist.

Let me run through an equation or two (I will use these a little later). If I allow for the energy of vibration in a particular mode, the time rate of change of that energy is going to depend upon the energy the mode would have at the local translational temperature. The difference between the energy the mode has at a particular moment and the energy it would like to have gives this type of equation, referred to as a relaxation equation.

It is the same equation you used as an undergraduate when treating radioactive decay. There is no difference. It is the decay from an excited state to a ground state. In this case, this tau, this relaxation time, is a measure of the time it takes a system to return to $1/e^{\text{th}}$ of the equilibrium state. Tau is related to some quantities we will refer to as rate constants.

A rate constant is a measure of the speed at which you get transitions from one internal state to another. It is for example, the number of transitions from the μ equals 1 vibration state to the μ equals 2 vibrational state per unit time at a particular pressure. In equilibrium as many collisions go up in quantum number as go down; that is what we mean by equilibrium. There is no change. This is referred to in chemistry as detailed balance. In a chemical reaction, the reaction is balanced when there are as many reactions in one direction as there are in the other. The same happens in these energy-transfer type processes.

Relaxation is entered into sound propagation through an effective specific heat. Going back to the old theory of communism, everyone shares equally. When computing the specific heat of a gas every degree of freedom gives one R of specific heat. For translation there are three translational degrees of freedom, so the specific heat associated with translation is $3/2R$. A polyatomic molecule has three rotational degrees of freedom which gives $3/2R$ more specific heat. This gives a total of $6/2R$ for the specific heat. Now, if you include vibration, you get more, maybe $7-8-9-10/2R$, associated with vibration.

You can identify a contribution to the specific heat associated with each internal mode. Specific heat is a measure of how much energy is stored when the temperature is increased. Energy can be stored in vibration, rotation, or translation. I can write the specific heat in terms of the specific heat of all those modes that are not relaxing plus a partial contribution from a relaxing mode. For example, if I am concerned with vibrational energy transfer, I can write the specific heat as a sum of the specific heat associated with rotation and translation plus a specific heat associated with the relaxing mode, the vibrational mode. The specific heat associated with the relaxing mode is denoted by C' .

Sometimes there might be many of these C 's, depending on how complex the molecule is. The transparency shows how we get from relaxing specific heat to a complex propagation constant. We can write the relaxation times in terms of these rate constants. For α lambda, we get the form that Dr. Garrett showed earlier.

We also find the dispersion that Dr. Garrett wrote earlier. I think I have it written as the inverse on the transparency. But if I invert, the speed of sound goes up with frequency. That amount that it goes up with frequency depends upon the specific heat of the relaxing degree of freedom. The amount of attenuation per unit distance also depends upon the specific heat of the relaxing degree of freedom.

There are other ways to put relaxation processes in the acoustic equations to compute the attenuation. The approach followed here is not unique, but it is one that is followed in most textbooks.

In the case of multiple relaxation, we have to be a little bit more careful. If I have a molecule, let us say sulfur dioxide -- physicists hate any molecule with more than two atoms and,

in fact, we hate most molecules with two atoms, except for hydrogen (hydrogen you can work with) – there are many vibrational degrees of freedom. During collisions, these vibrational modes are coupled. The result is relaxation of multiple modes simultaneously. You can no longer easily identify specific heat with a specific vibrational mode. You must, instead, consider normal relaxation pathways, which makes the whole problem a little more difficult but not a great deal. If you ever have a need to treat multiple relaxation, there are texts you can go to. The mathematics can become very tedious and messy, but there is nothing in the physics that we have not already talked about.

(Transparency)

I am going to summarize this first hour, and then we will take a break. Up to this point we have talked about viscosity, talked about thermal conduction, and internal relaxation. These are associated with momentum transfer, energy transfer, and the transfer of internal energy, respectively. All of these take place as a result of collisions. Collisions cause the sound wave to go from one point to another.

Would a sound wave go from one point to another without any collisions? That gets back to the question of can some waves propagate in a vacuum. I have already stuck my neck out on that one.

One topic that Dr. Garrett did not cover in great detail was relaxation absorption.

(Transparency)

As far as I am concerned, when we talk about acoustics that is where all the fun is so I am pleased that he left all the fun part to me.

(Transparency)

Studies of relaxation are very demanding from an experimental point of view. I will start this session as though I were Dr. Garrett and talk about experiments and what we can learn from experiments. Then we interpret those experiments.

I am interested measuring absorption or dispersion of sound as it propagates through a gas in order to learn something about internal energy transfer taking place at the microscopic level. There are some difficulties, as there always are, in devising a good experiment. The experiment should be simple. You get a speaker and a microphone. You turn on the speaker measure the change in arrival time and amplitude as the distance between them is changed. From that you get the attenuation and velocity. Simple.

(Transparency)

It turns out that the experiment is not quite that simple. In addition to relaxation losses, there are viscous and thermal losses in the gas that are proportional to f^2 . I would like for the relaxation absorption to be larger than the viscous or thermal conduction absorption, or at least equal to (if it is equal, maybe I can sort it out). If it is very small compared to viscous and

thermal conduction losses, I am probably not going to be able to accurately determine absorption and dispersion due to relaxation. Relaxation absorption and dispersion go up with frequency then flatten out; they do not continue to increase.

I want to measure the absorption at frequencies near the peak in the α lambda versus frequency curve. If I avoid high frequencies, I should be able to measure the relaxation absorption without too much contamination from classical losses.

Quite often, if I am going to measure absorption, I have to put stuff in a container. A convenient container is a cylindrical tube. If we constrain the sound to a cylindrical tube, there will be losses at the walls due to viscous drag. Dr. Swift makes a living out of this effect at the walls of the tube, but for us it is a hassle. We would like to minimize that effect. The viscous and thermal losses at the walls of the tube go as the frequency to the one half power. The associated absorption does not change very fast with frequency. We can, perhaps sort out wall losses by using this fact.

Another thing to note is that the losses at the wall diminish as the diameter of the tube gets larger. In order to minimize losses at the tube walls, you make the tube larger and larger. There is a limit however. If the frequency is so high that modes in addition to the plane wave mode are allowed to propagate in the tube, modal interference can make the received signal difficult to interpret.

Radiation leakage, which is greatest at low frequencies can also cause problems. If the tube is not very massive at the ends, the ends of the tube will vibrate and radiate sound into the surrounding space. Some of the sound is being lost not by absorption but instead is leaking out. It is hard to correct for radiation losses. The way to overcome this problem is to put thicker end plates on the tube.

So now you have a big diameter tube to reduce wall losses and you have very thick end plates to minimize radiation losses. We also want to operate at low frequency which means our tube must be long. Surely there is a better way.

One option might be to place the source in a large room and measure attenuation before the sound hits the walls. In that case, though, spreading losses become important. If I tried to do that measurement in this room, put a sound source here and a microphone back there, not only will absorption affect the received level but also spreading. Even for a directional sound source there will be some spreading.

At 1000 Hz absorption in this room will be on the order of one dB every 3 or 4 meters. A good measurement requires a 3 or 4 dB loss, so you would need a 20 meter propagation to get a decent measurement at 1000 Hz. In these 20 meters of propagation, unless you start out very far from the source, to begin with, you will have 15 or 20 dB of loss due to spreading. You can compensate for spreading losses, provided there are no surfaces around. So you must be several

tens of meters from the source and then have an additional 20 meters between you and the ceiling, floor and side walls. A very large room would be required. Either that or you use a tube that is large in diameter, has these heavy end plates, and is very long. Given the alternatives, the long tube does not look so bad.

The point I am making is that these measurements which seem simple can be quite challenging. Such measurements are especially tough in liquids. In order to make a reasonable measurement of the absorption of sound in air at 1000 Hz, I need a tube about 20 meters long and maybe a foot or two in diameter. Attenuation in a fluid, sea water for example, is about four orders of magnitude less. So, in order to make the same measurement in a fluid I would need a four-orders-of-magnitude longer tube; 20,000 km. I do not know about your labs, but ours are too small to accommodate tubes that long.

The measurement of relaxation absorption, although simple in principle, is often just not practical.

(Transparency)

There are some cases where one can make these measurements. This is one experimental apparatus that was used to measure sound absorption in a number of gases. I use it as an example not because it is very elegant but because it worked very well. What you see here is a transmitter and a receiver. The transmitter inside the gas chamber moves to increase the path length.

The sound amplitude goes as $e^{-\alpha X}$. If I measure the amplitude at several transmitter receiver distances, then draw a straight line through the log of the amplitude, I can determine α , the attenuation coefficient. Not only that, but if I transmit a signal continuously and look at the phase of the received signal -- referenced to the source phase -- then I can watch the phase change as the propagation path length changes. I can determine from that the velocity.

This was, by the way, a setup being used in Doug Shields' lab when I went to Ole Miss.

(Transparency)

Since then we have built a system which uses modern digital electronics. The idea is similar. There is a sound source and a receiver. In this apparatus we used tone bursts. We transmit a burst four or five cycles. Those four or five cycles propagate down the tube, which in this case was about 10 meters long, bounce off the end, come back, bounce off the end, etc.. We measure the amplitude of successive reflections and, from that, determine the attenuation coefficient.

Another device that can be used to measure the relaxation time is referred to as a spectrophone.

(Transparency)

Let us imagine for a moment that I have a test cell with a microphone in the side of it. Now supposed that I introduce through a window radiation from a pulsed laser that has the

proper wavelength, such that one of the internal modes absorbs the radiation. Molecules will absorb that radiation in a vibrational mode dictated by the wavelength of radiation. That vibrational mode will be temporarily overpopulated. But, due to collisions, the number of molecules in the excited state will decay with time.

As the number of molecules in the excited state returns to equilibrium, the energy that was placed in vibration appears in translation. The molecule temperature increases. The translational motion increases. If occurs, then the pressure increases. The transducer in the cell wall will see a pressure increase. The signal is an acoustic pulse resulting from the optical pulse.

Some interesting things happen, depending on the mode that is excited. For example, consider carbon dioxide. Carbon dioxide has multiple energy levels associated with bending of the molecule, symmetric stretch, and asymmetric stretch. An interesting feature of CO_2 . If you zap CO_2 with a CO_2 laser operating at 10.6 microns, molecules are promoted from the symmetric stretch which is closely coupled to the bend to the asymmetric stretch. The symmetric stretch is left under-populated. The first thing that happens is that this under-population is replenished from the bending mode and translation. This happens before the asymmetric stretch can transfer excess energy to translation. The first thing that happens to the gas is that it cools; it does not heat.

Think about that for a refrigerator. Where is Dr. Swift?

(Laughter)

DR. MIGLIORI: Run that by me, again?

DR. BASS: It is an interesting phenomenon.

DR. MIGLIORI: Okay, you drive the asymmetric stretch, which is the IR-active mode, because you are hitting it right on the money, so it is heavily populated. I have other modes. So if I start to overpopulate, the first thing that overpopulates is the symmetric stretch?

DR. BASS: Yes.

DR. MIGLIORI: So now that is heavily overpopulated and so I need a translational mode and the bending mode to get it into a stretch or something so that it starts to suck energy out of that. I see.

MR. POESE: And that works only because the mode you have overpopulated is a higher energy mode than the rest of them.

DR. BASS: That is right. It would not work if the mode you overpopulated was efficiently connected directly to translation. It happens only if you have something like a metastable mode, which, of course, you have for the laser system, to begin with.

DR. MIGLIORI: And it needs the transitional energy to strongly couple to the next type of mode. There is some selection rule at work.

DR. BASS: In molecular lasers, one of the requirements is that the lasing vibrational mode be metastable, which means that most gases that lasers are going to behave like this.

I have described ways to observe relaxation processes and measure relaxation times. Let me discuss a few results.

(Transparency)

This transparency shows $\alpha \lambda$ for fluorine. Keep in mind this curve is α times λ . If we plotted the absorption coefficient, it would go up here and flatten out. You might ask yourselves why in God's name would any work with fluorine? Fluorine is one of the nastiest gases known to man. It dissolves most everything, including glass.

The reason is that fluorine has good absorption of sound; it absorbs fairly strongly. An apparatus of reasonable size forces one to select gases that have vibrational relaxation absorption that is large; fluorine is one of those gases. Doug Shields actually made these measurements; I would not dare. Not only that, he made measurements with the fluorine was heated, which makes it even worse.

DR. MIGLIORI: Yes, but the key thing is says kc/atm, which means it was done before the EPA got on everybody. You could never do that today.

DR. BASS: That is right, you could never do that now. I am not sure we would even be allowed to have fluorine in our labs any more.

When making these measurements you can increase the frequency or you can decrease the pressure. The two effects are equivalent, because the important thing is the ratio of the mean free path and the wavelength. That ratio can be changed by decreasing the pressure or by increasing the frequency.

If you look on the transparency, you will see different types of points. These points represent measurements made at two different frequencies, but the pressure is adjusted so that the ratio of frequency and pressure are the same. Note that the result is the same. So $\alpha \lambda$ really is a function of frequency divided by pressure.

The other curve shows velocity dispersion both measured and computed. The velocity, continues to increase as the frequency increases and different types of relaxation processes are encountered. Rotational relaxation would cause the velocity to go through another step. Translational relaxation would result in another increase. Note that the data fits the curve pretty well.

(Transparency)

This transparency shows more experimental results. These results are for chlorine. Chlorine is almost as nasty as fluorine; not quite as bad, at least it does not dissolve glass. Results are shown for different temperatures. As you can see, the magnitude of the attenuation increases with temperature and the relaxation frequency increases with temperature.

(Transparency)

The prior results are summarized on this graph which shows the probability of a vibrational quantum being exchanged as a function of temperature. On this graph low temperature is to the right, high temperature to the left. As you can see, the probability of energy transfer during a collision decreases as the temperature decreases. The probability is greater for iodine than it is for chlorine, which is greater than it is for bromine.

(Transparency)

Now to draw some conclusions. First, the log of the probability goes as the temperature $^{-1/3}$ power. The explanation for that dependence is simple. The vibrational quantum levels that are separated by about 800° K in temperature units. Not many colliding molecules, but quite a few, will have velocities with sufficient energy to transfer a quantum from one molecule to the next. You are out on the tail of the Boltzmann distribution, but not far out.

As the temperature increases, more and more collisions will involve molecules that have velocities great enough to cause a vibrational transition. That is all this temperature dependence really means. This temperature dependence can be predicted by a quantum mechanical scattering calculation. The derivation is complex but the underlying physics is simple. Except for the one-third, you could probably deduce a similar temperature dependence with just a hand-waving argument.

The probability of energy transfer during a collision is very sensitive to the energy deficit for the same reason. If two molecules collide and they do not bring into the collision sufficient energy to make a quantum of vibration, then there will probably not be an energy transfer (unless you get quantum mechanical tunneling - even then it will be a very, very low probability event).

The probability of energy transfer also depends on how hard the collision is. When molecules collide they interact according to some interaction potential. Often when two molecules begin to collide there is an attraction and they speed up. But then when they get really close they begin to see the nucleus, the hard core, of the other molecule. When that happens, they begin to repel. So there is a speeding up then a repulsion.

The probability of energy transfer depends on the nature of the interaction potential. If two molecules start out with a large attraction, they will accelerate towards each other until the hard core is reached. At that point they will be going with much greater speed than they started out with. If energy is needed to make a quantum of vibrational energy they might now have it.

On the other hand, if they interact via a shallow potential, the transition might not occur. The probability is going to depend on how hard the collision is, how fast the collision takes place.

The probability of energy transfer also depends on the geometry of the mode. I think that is fairly easy to see. Consider CO_2 . CO_2 has three modes. Two of them are like this and one of

them is like this. If you were sitting there throwing rocks at a CO_2 molecule, which would be easier, to hit somewhere along here and cause the bending mode to change, or to hit directly on the end and cause symmetric stretch to change? Obviously, if you are throwing rocks, the bending mode has a much bigger cross-section than does a stretch, so you are more likely to excite the bend. It is a geometry effect.

(Transparency)

This transparency shows an interaction potential. I have already talked about how the potential affects the energy transfer probability.

I am going to take about one more example molecular system: sulfur dioxide and oxygen. Sulfur dioxide and oxygen comprise a very complex system. Oxygen has one vibrational mode. Sulfur dioxide has three vibrational modes. When an oxygen that has vibrational excitation collides with SO_2 , what is going to happen? Is the oxygen's vibrational quanta going to go to translation? Will it go to SO_2 vibration? If it goes to SO_2 vibration, which vibrational mode?

We can use a few rules of thumb here. First of all, the oxygen vibration has an effective temperature of about 2400° , so the probability of transferring all of the energy to translation is almost zero.

SO_2 has three modes. One of those modes is almost resonant vibration with the oxygen vibration. The rule of thumb is you always take a vibrational relaxation pathway that minimizes at each step the energy deficit. So we would expect, then, the energy to transfer from oxygen to SO_2 and then cascade down through the SO_2 internal vibration energy levels to translation.

Experimentally mapping the energy transfer pathway is not trivial. Plotting the pathway requires measurements at different concentrations. A number of such measurements are shown in the following transparencies. I will not go through all of them, because they are in your books.

(Transparency)

The final result is that the energy transfer appears to take place in sulfur dioxide as shown on this transparency. If oxygen is excited, there is rapid v-v coupling to this mode. This mode in sulfur dioxide is v-v coupled to that mode and you get cascading down.

(Transparency)

Another interesting example of vibrational energy transfer is carbon disulfide. The relaxation of carbon disulfide is interesting in that carbon disulfide can easily be handled in the lab both as a liquid and as a gas. It turns out that the two states relax very similarly. In fact, you can treat liquid carbon disulfide as a very dense gas.

Liquids typically relax much faster than gases. The molecules are packed more tightly together, so there are more collisions per second, about 10^{12} , rather than 10^{10} for a gas. The molecules in liquids are packed together so tight that quite often there are loose bonds between

them. So there is no longer one molecule interacting with another, but clusters of molecules interacting with clusters of molecules. Now, clusters of molecules have lots of energy states available to them that individual molecules do not have, so there is a much richer array of energy states for energy transfer to take place through.

As a result, one would expect relaxation to be much faster, more than two orders of magnitude faster, in a liquid state as compared to a gas state. But in sulfur dioxide that is not true. We can treat liquid carbon disulfide as a very high-pressure gas.

That is not true for H_2O . The reason is that water forms those clusters I mentioned. The reason for the difference is that H_2O has a significant dipole moment. Carbon disulfide does not have nearly so large a dipole moment.

One application of vibrational energy transfer has to do with absorption of sound in air. Our atmosphere is composed of primarily nitrogen and oxygen, with a little dash of CO_2 and a little sprinkling of water (at least in Monterey - in Mississippi we have a large sprinkling of water), and just a little bit of argon thrown in for good measure. The propagation of sound through the atmosphere has historically been a problem of interest. It is something that affects us all on a daily basis and has received the interest of scientists over the past century and a half.

(Transparency)

In considering the absorption of sound in the atmosphere, we can compute the classical and the rotational relaxation absorption. I am not going to go through the details, they are given in the handouts. The bottom line is that for frequencies of communications interest, losses due to viscosity and thermal conduction can be ignored.

The only losses that are important are those losses due to vibrational relaxation. The only thing that can keep me from being heard 20 miles away on a quiet day are quantum effects in molecular collisions. If it was not for quantum mechanics, you could hear me at much greater distances. In order account for all relaxations which influence sound propagation, I have to include interactions of oxygen and nitrogen, oxygen and oxygen, CO_2 and oxygen, oxygen and H_2O , etc.. Each of these contributes to the vibrational relaxation of oxygen and nitrogen which have most of the relaxing specific heat.

It took five or six research groups about 15 years to measure all the rates which are important in air. But after measuring all these rate constants, we can put them into the relaxation equation I showed you earlier and predict absorption of sound as a function of frequency. The result is shown on the following transparency.

(Transparency)

The dots are experimental measurements. The line is prediction. I would argue that that agreement is fairly impressive. Hopefully this level of agreement convinces you that the theory I have been talking about has some merit. May

This problem, which has been worked on for more than a century, is now closed. We know what causes the absorption of sound in air. There are no longer any mysteries in that regard.

DR. MIGLIORI: How come the data say 266° K and not 300 or 500?

DR. BASS: Actually, we went through a series measuring absorption every 10°. The steps had to be Fahrenheit, because we were doing it for NASA. That is why you get the peculiar units.

(Laughter)

DR. MIGLIORI: But you would expect to have trouble at high temperatures in air. No?

DR. BASS: We took data up to about 110° F. We did measurements from 5 percent relative humidity, which is all we could get consistently in Mississippi (actually, we had bottled air), to 98-99 percent relative humidity at each temperature. The graph you see here is typical. We tried to find some reason to make more measurements. It is hard to justify.

(Transparency)

In later lectures you will hear about sonic booms. In atmospheric acoustics one of the interesting effects is that very high-amplitude waves evolve into N-waves due to nonlinear effects. The absorption mechanisms that I just described cause these N-waves to try to become sine waves.

So there is a constant competition between steepening and rounding. Rounding because of absorption due to relaxation processes and steepening due to nonlinear effects. The final shape of a sonic boom depends upon the amplitude.

(Transparency)

The last topic I want to talk about in any detail is absorption due to other processes. The first of those is diffusion. I should warn you that this part of my lecture will be less rigorous. I started out talking about a topic I know a lot about then moved on to subjects I know something about. Not much, but probably more than you. I am now going to delve into topics I know very little about. Most of you will know less, so maybe we can learn together.

DR. BROWN: I have a question before you go on. You skipped over a viewgraph before the shock waves that shows the two prominent relaxation phenomena, nitrogen, oxygen, and total loss. There is no scale under frequency.

At 10 kHz is nitrogen relaxation the dominant effect?

DR. BASS: The answer depends somewhat on the relative humidity; that is the reason there is not a frequency scale. If I were to consider the air in this room right now, the peak for nitrogen would probably be at 50 or 60 Hz. The peak for oxygen would probably be at 1000Hz or so. In Mississippi, the nitrogen peak might shift to 1 kHz, and the oxygen peak would also shift up. That exact shape depends a lot on humidity.

DR. BROWN: I am really talking about the curve there, where you really have a transition.

DR. BASS: Where this break occurs depends on humidity. The break due to oxygen is right here. That also depends strongly on humidity. These break points for oxygen and nitrogen are very humidity-dependent, that is why there is no frequency scale.

DR. BROWN: I do not see a second break point, a second inflection point, on the curve on the prior viewgraph.

DR. BASS: On the curve where I had data? That curve extends to 10 kHz. Oxygen relaxation dominates. It is only at high temperature, high relative humidity, that the nitrogen relaxation is shifted over where you can see the second deflection.

DR. MIGLIORI: But as you increase the humidity, that whole nitrogen curve moves up so that it intersects at a different spot. That is, all the attenuation goes up, right, everywhere?

DR. BASS: As you increase the humidity, the primary effect is a shift to the right, you shift everything to the right.

DR. MIGLIORI: And they are higher.

DR. BASS: Slightly higher.

DR. MIGLIORI: So does the attenuation decrease at low frequencies as you increase the humidity?

DR. BASS: It decreases at low frequencies as we increase the humidity.

DR. MIGLIORI: But at higher frequencies it goes up?

DR. BASS: At higher frequencies it goes up so long as you are below the relaxation frequency.

(Transparency)

One of the contributions to absorption of sound in gases that I have not talked about up to this point is due to diffusion. I think Dr. Garrett mentioned this briefly. Consider a gas mixture of two atoms with much different masses, maybe xenon and helium. As the sound wave moves through, these two different gases are going to follow changes in temperature at slightly different efficiencies. Collisions do not take place as rapidly in the heavier gas as they do in the lighter gas. The result is some phase difference between the responses of these two species. That phase difference will lead to absorption. The form of the absorption is shown on the transparency.

For most cases this absorption is not very important. In air, diffusion absorption can be shown to be μ less than 0.1 percent of the attenuation due to viscosity, which itself is only about 1 percent of that due to relaxation processes.

But there are some cases where this absorption is important. One case is absorption of sound in uranium hexafluoride-helium mixtures. In such mixtures, this term can become quite large. This is a topic that has been explored in only a few cases. There was a graduate student at

Colorado State University 10 or 15 years ago, who made measurements helium/xenon mixtures. His measured absorption did not follow this equation very well.

It is not clear why. We are not sure how sound propagates in mixtures of species that have such different masses.

Some of you probably are familiar with propagation of sound through porous media.

If not, think of sand, sound propagates through the pores. It propagates relative to the frame. Assume that the frame becomes elastic, so as the sound propagates through the frame, the frame also moves. The pore fluid and frame move relative to each other and not with the same speed. We often talk about there being a slow wave in the system and a fast wave in the system.

Assume we disconnect the connections between all the sand grains. Now we have sediment, mud, silt. Sound propagates in the fluid and the fluid interacts with the silt, so still there are still two waves. This situation was described by Biot. Waves exist in the pores and waves exist in the matrix and they move relative to each other.

A similar type of effect should take place in a gas but it has never been observed.

(Transparency)

Another interesting problem relates to a question Dr. Hofler raised earlier.

Consider an acoustic wave propagating through a gas. When the acoustic wave enters a parcel of gas it disturbs the local equilibrium. In the absence of a disturbance, there is some equilibrium distribution of molecules amongst translational, rotational, and vibrational energy states. If left alone, through collisional processes that distribution will remain constant.

But the sound wave causes the translational motion to change, so all these other modes have to change as well. The system leaves one distribution and tries to get to a new distribution which would represent equilibrium with the new energy available.

The only way to go from one distribution to another is through collisions. In equilibrium there are as many molecules going from one state to the next as there are going back, but in the presence of a sound wave that is no longer true. This distribution function evolves with time via collisional processes.

One can formulate this problem in terms of the rate of change of the distribution function, the function which describes how many molecules are in each state. The rate of change of that distribution function can be determined by solving a collision integral, provided the perturbations are small and the interaction between the colliding molecules is known.

Refer to your handout for more details.

The Boltzmann equation allows you to describe the transition from one statistical distribution to another.

Assume I have a gas at very low pressure in a container. In the wall of the container there is a transducer that disturbs that gas, that vibrates. Further assume that I receive that disturbance at some other point.

If I lower the pressure enough, I reach the point where a molecule leaving the transmitter can get to the receiver without having a collision. Molecules that are moving more slowly will have a higher probability of suffering a collision and they may be scattered off in another direction. Molecules that are moving very slowly will probably suffer many collisions and never transfer their information to the receiver. As a result, the speed of sound that one measures will depend upon the distance of propagation. The greater the propagation distance, the more likely it is that there is going to be a collision and the most likely molecules to participate in that collision are lower velocity molecules.

The propagation of sound through the gas is characterized by a propagation constant that includes absorption, which is proportional to distance. The term that goes into the $e(\alpha X)$, the α itself depends upon distance. The velocity also depends on propagation distance. It goes as the distance squared. So you have a velocity which depends upon propagation distance and absorption which depends upon propagation distance because only the high-velocity tail of the distribution is getting to the receiver.

This result has been predicted by several people. Boltzmann might have written down some of this work initially, but there have been other, more accurate, representations of this equation within the past 30 years. But this distance dependent absorption and velocity has never been observed. The reason there has never been an experiment no such measurements have been made since the advent of modern electronics. Another opportunity for a student seeking a Ph.D.. The technology and the instrumentation are available to do these measurements but they have not been done.

(Transparency)

Such measurements would allow one to learn about the motion of individual molecules. That is an entirely new game; an unexplored area. It is pretty easy to describe the transition of equilibrium distribution A, to equilibrium distribution B if you can use statistics to connect the two.

It is a tough problem when you have in your box 48 particles. There are not enough molecules assume statistical distributions any more, but there are many too many to try to solve for all the collisions that are taking place. That is a realm that current physics has not been able to effectively deal with.

Similar problems crop up in disparate fields. Dr. Hargrove was talking about some work that his son is doing. He is making circuits so small that the carriers can no longer be treated

statistically. They have to try to account for all the individual carriers, but there are too many carriers to account for individually.

MR. PRATHER: What do you mean by large distance? What I mean by the question is, if the distance gets larger and larger, then even the fast molecules stand a chance of colliding, but as the distance increases, of course, the slow molecules have a greater chance of colliding.

DR. BASS: Let me give you some ballpark numbers. In helium at a frequency divided by pressure of about 500 MHz/atm, the distances required to be long are on the order of 1-2 mm.

MR. PRATHER: But you do not want to keep getting larger and larger, right?

DR. BASS: You want to get larger and larger until you lose your signal entirely. As the propagation path gets larger, these effects become more pronounced. But after a while you will get no signal. There is a limit to how far you can go even with modern electronics.

There are a few other topics I considered talking about. Some of those are outlined in the handout. Given the time of day and how long I have been talking, I am not as eager to cover some of these topics as I was when I made out my notes. So I am going to stop at this point and give you a chance to ask questions.

STUDENT: In your summary what energy is being exchanged.

DR. BASS: In the case of standard thermal conduction, the quantity being transported is translational energy. When internal energy modes are present, quantized states of vibration and rotation are also transported. In both cases you are talking about transporting energy; the type of energy that you are transporting is the only difference.

DR. HALLEJ: For transmitting sound waves, say a plane wave, the bottom line is propagating translational energy? The rotation of the molecules seems not to be as important now, right?

DR. BASS: You cannot separate rotation from translation. As the sound wave propagates through the gas, each mode is going to be excited and have its affect on propagation. But your receiving transducer is going to be sensitive only to the translational motion.

DR. GARRETT: Isn't one way to respond to that to go back to the Laplacian sound speed, the $\gamma RT/M$? As you start piling on these other degrees of freedom, you start dropping that γ way down.

DR. BASS: Good point.

DR. GARRETT: It manifests itself in that reduction of speed, because you are putting energy into the non-translational degrees of freedom.

DR. BASS: If you had a polyatomic molecule that had several degrees of freedom, γ might be 1.1. When those vibrational modes freeze out, γ is about 1.4. At higher frequencies the rotational modes freeze out and γ is near 1.7.

DR. GARRETT: Yes, that is right. As the molecule gets more complicated it shows up in that gamma. In fact, had people been aware of that and measured the specific heat of hydrogen as a function of temperature, quantum mechanics would have been discovered 20 years earlier, because at low temperatures the specific heat is $3/2$'s at constant volume, because everything is frozen out.

Then you get up to about 100°K and the rotations are in but the vibrations are out. You get to about 2000 to 3000°K , then the vibrations are excited, so you see the sound speed or you see the specific heat changing with temperature in that way, but quantum mechanics was not known at that time and nobody picked up what was actually a quantum mechanical effect that was perturbing either their sound speed or specific heat measurements.

DR. MIGLIORI: I do not understand what it means to say that the faster molecules are responsible for propagation. It seems to me that as you get out of this sort of near-field region, however you want to describe it, where you have problems with a fast molecule going a long ways and a slow one going a short way, so that there is some spreading out of the sound pulse on a length scale of the order of the mean free path, but once you have gotten beyond that, the whole distribution function is not very much perturbed, right?

DR. BASS: In the conditions where you have very low pressure, you never get beyond that region. The process is probably better described as a diffusion wave and the distribution just gets spread out over more space the more you propagate.

DR. MIGLIORI: So the mean free path tells you how sharp a feature you can propagate cleanly?

DR. BASS: Yes. The numbers I gave a little earlier at those F/P 's, were the mean free paths. That is where the wavelength and the mean free path are essentially equal.

DR. MIGLIORI: At distances where the propagation path and the wavelength are much greater than the mean free path, I am back into some sort of stable propagation regime, again. The distance over which I propagate is large, compared to the mean free path, and the wavelength is large, and all I have to do then is hope that that regime happens before other attenuation processes kill it?

DR. BASS: But if your wavelength is much greater than your mean free path, you are going to get fairly low attenuation. If the wavelength be much longer than the free mean path the attenuation is small. High attenuation and spreading of the distribution, by definition, occur in the other extreme.

DR. ATCHLEY: What does the wavelength mean in a regime where molecules travel differently going across the tube?

DR. BASS: What does wavelength mean when you have a velocity that is not only dispersive but also varies with distance. I am not sure I have a good answer to that question, Anthony.

DR. HAMILTON: When you have a transient, for example, in a viscous fluid -- just take a step, a pressure jump -- as that propagates, the foot of that step will propagate out from the wave. If you track a certain amplitude or wavelength, for example, depending on which part you track, you can get an arbitrarily fast speed associated with it.

On the other hand, if you forget about viscosity and use the equations from this morning and consider a lossless fluid, presumably Dalabere's solution holds. Then you have an abrupt jump for, for example, a single cycle of the sinusoid. That zero will move just at C_0 . How does that fit in with how you discussed diffusion. I have always had trouble explaining that when we talk about these dissipative processes in class.

DR. BASS: In the ideal case where you do not have any frequency dependent losses you are allowed to propagate one single burst of a sine wave or any other wave shape unless you include nonlinear effects. If there is a jump in pressure and the media is dispersive you must consider that jump in terms of the Fourier decomposition; there are high frequencies and low frequencies present.

The high frequencies will be attenuated, so you will no longer have a leading jump, a square jump. Instead, you have this shoulder. That is the way I explain it.

DR. HAMILTON: That is the way I describe it, as a filtering effect. But it always struck me that I was hand waving. Mathematically it makes sense. But physically, if I hit a fluid with a transient or a piston where do those components come from. At $T = 0$ there was no wave out there, there was no disturbance to come out from the ashes, way out halfway to infinity, or something. But the equations will tell you there is a disturbance out there. When you solve the equations, the transient will move arbitrarily fast ahead once you put frequency dependent absorption in. Or am I not correct?

DR. BASS: It is a matter of magnitude, but I think that what you say is correct. I am not sure how to respond. I suppose that in any given situation there is going to be some very fast molecules, and those very fast molecules are going to transmit information more rapidly than the slower ones.

Whether or not it goes out to infinity is a different question. I do not know the answer.

DR. MIGLIORI: That is sort of where the diffusion constant came into this stuff. You have this pressure pulse and it is propagating and it is diffusing.

DR. BASS: I am not sure that is properly called diffusion, but I will --

DR. MIGLIORI: I do not understand this completely. I mean, it is obvious that one of those molecules is going to make it to the end of the propagation path at a speed that is

determined by its kinetic energy when it started, which could be a very high number. It could be on the upper side of the distribution function.

If you have 10^{23} molecules, one of them may make it without a single collision, even though it is many, many mean free paths. It is the tail end of a whole bunch of statistical distributions that does it, and so you would think that somehow diffusion plays a role in that somewhere.

DR. BASS: I do not think that the diffusion equation that we normally write will give you the proper form for the solution.

DR. GARRETT: I think you might be able to clarify this problem of propagation -- basically, you are stating Dr. Atchley's comment. He asked what is the meaning of wavelength? I would like to put that in another way, which is that when we talk about the transmission of information by waves, the thing that is unique is that the material does not move from point to point; the energy moves from point to point.

That is, you put a cork on the ocean, you see the waves coming in to the shore but the cork just goes up and down. When you go to the limit you are talking about, I would say that that is ballistic, it is not an acoustic limit. Wavelength does not make any difference, and you are not doing waves any more; you are shooting particles from this location to this location and, therefore, it is not really a wave. It is artillery.

DR. BASS: Thank you.

(Laughter)

DR. GARRETT: Professor Bass' old occupation -- I am sorry, it is lieutenant colonel, Col. Bass' former occupation.

DR. HARGROVE: What happens if we try to build a resonator that is only a few half-wavelengths, whatever wavelength means in this rarefied medium regime? Is there any new physics to be learned in trying to think of a resonator in the rarefied regime?

DR. ATCHLEY: Where is the node and the anti-node? What is the variant?

DR. BASS: There have been a number of computer simulations and they appear in the literature from time to time.

STUDENT: Maybe molecular dynamics?

DR. BASS: They appear primarily from people who work in statistical mechanics. The last one I have seen was some work by Seth Putterman. There has not been, to my knowledge, an experiment that has been proposed or carried out.

DR. ATCHLEY: Are the collision processes well-enough understood to do this experiment? How do you get the outliers? How often do you have to wait to have an outlier? Because when it hits a wall, it is going to bounce off with an energy determined by the wall among other things.

DR. BASS: I think that you would want any wall other than the signal transducer to be several times further away than the receiver. In the regime we are talking about, propagation path lengths in millimeters are not unreasonable. I do not think it will be that difficult to avoid walls other than the transducers themselves.

DR. HOFER: I have a question that is related. You talked some about hydrogen and rotational excitation. I have looked at data books describing ortho-hydrogen, para-hydrogen and normal hydrogen. I have never resolved this stuff myself. Perhaps you could list these things and what they mean in terms of real measurable or phenomenological properties.

DR. BASS: I suspect there are other people here who can speak to that better than I can. Hydrogen molecules, for reasons that I now forget, have even rotational quantum numbers or odd rotational quantum numbers, and you cannot have transitions between odd and even if the molecules are free (that only happens during interactions between molecules).

A molecule can go from the j equal to 0 state to the j equal to 2 state, or j equal to 4 state, but it cannot go from $j=0$ to $j=1$. It can go from 1 to 3 but cannot go from 1 to 2, for example. Molecules that have only odd-numbered rotational quantum numbers are called, I think, para -- I say that because it is just opposite from what you would think. The nomenclature comes from deuterium which is just the opposite.

The molecules that have even numbered rotational states are referred to as ortho. So you have ortho- and para-hydrogen because of the difference in the allowed rotational states.

During a collision, instead of having to go from, say, j equal to 1 to j equal to 2 to cause re-equilibrium in hydrogen, you have to go from j equal to 1 to j equal to 3. This is one reason why hydrogen rotational energy transfer is fairly slow. The rotational energy level gaps are effectively larger.

You asked what normal is. Because of the number of different ways of achieving para-hydrogen and achieving ortho-hydrogen, they should be in a standard ratio (the statistics are about 2 to 1, I think), so normal hydrogen would be a 2-to-1 mixture of para and ortho.

Those of you who know something about statistical mechanics, have I screwed that up really bad?

DR. MIGLIORI: Isn't it 3 to 1? When you are in a para state there are three quantum numbers and, in the ortho, it is one? And there is some very long time constant when you liquefy hydrogen -- I have forgotten what it is -- for it to stabilize.

DR. BASS: Para and ortho are fairly long-lived.

DR. MIGLIORI: Yes, like many hours. Maybe it is days. And it is the nuclear spins, right, that align? So when they are parallel, you have three quantum numbers, three states, both up, both down, and the in-phase contribution of the two states, and, when it is ortho, you have

only one. So if you equipartition all those, you get three and one, so it is 25 percent. I am playing this by ear here.

DR. BASS: I will take your word for it.

DR. GARRETT: The slow relaxation processes that you are quoting would take 10^{11} collisions. The only way you can appreciate that is to realize that there are only about 10^9 people on the planet. You would have to meet the right guy 100 times to have the same probability that one of those collisions is going to give you a vibrational transition.

The unlikelihood is so enormous, if there are 10^9 people and only one of those people only 1/100 of the time is also going to do you any good, you can see why those processes are basically irrelevant.

DR. HOFER: If you had a container of room-temperature hydrogen, would any of that matter? If you were doing some kind of measurement, would you see anything?

DR. BASS: It turns out that if you measure acoustic attenuation in the frequency range of standard measurements, the presence of para and ortho does have a subtle effect. As a result of their existence, the rotational relaxation time of hydrogen is much longer than it would otherwise be. At reasonable frequencies, then you will measure a relatively larger attenuation than you would if such states did not exist.

DR. MAYNARD: On these subjects, of all the books on my bookshelves, two of my favorite physics books are books on this subject. I do not think the handouts to the students included them. There is a book by Ter D. Haar, *Statistical Physics*. It is a beautiful book that explains how you go from microscopic collisions to these collisions.... It is a beautiful book. The other book, by Canard, is on kinetic theory that covers this regime between the microscopic and the macroscopic.

STUDENT: When you have different gases, how many waves propagate, two waves or one wave?

DR. BASS: In most cases that I am aware of you can always describe propagation by a single wave.

STUDENT: With effective mass?

DR. BASS: With effective mass. Of course you also have the possibility of a diffusion wave. There are alternative ways to look at the problem where you do get more than one propagating solution. To my knowledge, no one has observed more than one propagating solution in a mixture of gases.

DR. MIGLIORI: This is sort of an inertial separation. That is observed in solids with a charge separation, but it does not produce two separate wave speeds. It simply changes the restoring forces, because you have mucked up the distribution function in a way you would not

normally do it, that is, inertially separating the heavier and lighter species. It just changes the restoring forces and the wave speeds.

DR. BASS: But if you look at the propagation of sound through sand, you can see two distinctly different propagating waves. I am not sure that I can explain what distinguishes between the cases. Maybe Dr. Garrett is going to explain it for us.

DR. GARRETT: If you have an elastic matrix, then you can separate the two modes. You have a mode with the fluid passing through the pores and the pores are basically stationary, and you have a counterflow. That is why I do not think you will see it in a mixture of heavy and light gases.

As you go to the case of superfluids, you have, again, the two velocity fields and your fluid. But there are two different restoring forces. When you look at a diffusive system, you do not have a restoring force that involves the separation of the two media as you do in, say, in a superfluid, where the entropy difference is carried around in one of the fluids and, therefore, you have a restoring force that is thermal. You do not have that in the case of a heavy and a light gas.

You have it in the case of a rigid matrix in a fluid, because the matrix has the elastic restoring force and the fluid has compressive restoring forces. But I do not think you will see two modes in a gas. If you stretch the definition, you get a propagating mode and you get diffusion.

DR. MIGLIORI: It is not a force, right?

DR. GARRETT: Exactly. It is just diffusive. So if you are willing to stretch that, to say, well, you have a diffusive mode and you have a propagating mode, then maybe you have two modes. If you are talking about two propagating modes, I would not buy it, because you do not have a restoring force that is proportional to the separation as you do with, say, a superfluid, or as you do with a rigid matrix with a compressible fluid. That is the phenomenological picture.

RUS AND MATERIALS PHYSICS

ALBERT MIGLIORI
LOS ALAMOS NATIONAL LABORATORY

DR. MIGLIORI: Most of my viewgraphs are from other talks, dozens of other talks, and a few of them are new. Also, there will be a little bit of information I will show that I did not give you notes for, partly because I am in the midst of writing a book on this subject, with Wiley, and they kind of did not want to give away the store before you guys actually paid for the book.

(Transparency)

I get a little bit when you buy the book too, so you should remember that.

(Laughter)

I promise to make more mistakes than Steve Garrett did, and find fewer of them. I am going to tell you more than you ever wanted to know about resonant ultrasound, and I am going to tell you at a level where I guarantee to leave you in a position of getting into trouble without going into more detail on most of these points.

Whenever I think about this project, the whole thing is rather amusing. Part of the reason is that -- we will get into a little bit of a minimalist approach to physics before we dive into it -- measuring the resonances of a pinhead-size sample of high-temperature superconductivity on the lab bench, 200-year-old, or older, physics done very simply has gotten this huge crowd of people involved, eventually, at the end.

I have had wonderful students -- John Sarrao got his Ph.D. project on this. I have involved workers in all areas of Los Alamos from theorists to metallurgists, people working on nuclear weapons (which I will mention a little bit at the end), and we have also had strong interactions all over the country, all over the world.

But the weirdest thing is that as this project started to grow out of control at Los Alamos, suddenly we found hundreds of administrators who really needed to help you. Administrators, of course, never help you with anything, but they have this tremendous need to believe they are helping you and this costs you a lot of time and makes them feel better.

Well, once hundreds of administrators decided that this was somehow a success for Los Alamos, both our state Senators got involved and both of them will tell you how they supported this work and got it going at Los Alamos. Of course, once the Senators got involved, the T.V. stations got involved -- we were on television several times. At the very end I will show you almost the most ridiculous thing I have ever done in my life.

The Department of Energy finally realized that something good was happening at Los Alamos, so they got directly involved and, in fact, the hand-wired circuit boards for the resonant

ultrasonic electronics that I designed were shown to Congress at one point as an example of how technology transfer -- this is a buzz word -- works at Los Alamos.

Finally, the state of New Mexico realized something interesting was going on and they gave us a highway bridge. The bridge was six lanes of interstate over the Rio Grande, Interstate 25. The bridge was slated to be torn down and they gave it to us to do a destructive test to see if we could see how the resonances of a highway bridge change when you put in damage.

In fact, that was a circus. We had all of the T.V. stations going and we took this main 12-foot I beam on this bridge and we cut it with a cutting torch, starting from the middle of web. Of course, as soon as it went through the flange at the bottom, the bridge went clang. But we measured the resonances before and after and we actually had enough signal to noise to see a 2-inch crack in a 12-foot I beam on a 70-foot span of highway bridge.

All of this, of course, was in some sense, accidental; I never planned this to happen. But it did, so I am going to tell you how to do it, meaning resonant ultrasound, not this stuff.

(Laughter)

I will tell you what it does and that is where the really interesting stuff is. Of course, that is where you will fall asleep after the second break. Then I will tell you about nondestructive testing, which is where the money is and the aggravation is. Of course, in this particular case it has produced 17 new jobs in Albuquerque, which makes all sorts of people really happy that they kept cutting our budget but did not quite cut it to zero, because they can claim they funded us, but they never actually spent any money on it.

DR. SWIFT: Hundreds of administrators to create 17 new jobs.

(Laughter)

DR. MIGLIORI: Dr. Swift always goes to the heart of the matter, never misses the point.

(Transparency)

The first rule of life is to maximize the number of smiles, but the matrix is exceedingly complex. One of the things that always makes me smile is a minimalist approach to things and I wanted to give you a few examples before I show you why there was a crying need for somebody to take 200-year-old technology and indicate that it might actually be useful, again.

If you look at a crescent wrench, this is a very small problem and it is a problem that you can solve. The reason you can solve it is that there is only one basic design left at the hardware stores (with minor variants).

In contrast, something like chess is basically an insoluble problem in spite of computers. The reason is, it is just too complicated. The proof that it is not soluble is that people still play the game.

(Transparency)

You can carry this over to physics a little bit and you can look at particle accelerators. I am going to show you, at several levels, why resonant ultrasound is better at solid-state physics than particle accelerators.

These things are absolutely an insoluble problem. The proof is that no two particle accelerators are alike. There really is not a minimalist approach possible, I think, for something this complex. Great physics comes of it, but I do not think the physics is any greater than the physics that comes from something like the Josephson voltmeter, which is really benchtop physics, and that is what I love.

This is a wonderful macroscopic manifestation of a quantum physics problem. It is really cheap -- and of course today we define "cheap" as costing less than a car and smaller than a kitchen -- and it redefines the standard for voltage, among other things, just as good as the physics that comes from this thing and maybe even more important in some ways.

So the diagonalization of any research matrix, somehow the money and the size of the experiment just do not come into it.

(Transparency)

However, the particle accelerator often attacks very interesting physics problems that have to do with solid-state physics. The physics problem that you can put in the back of your minds for a while that has to do with particle accelerators is this. I am going to be talking about solids, and solids are really composed of little tiny balls. The little tiny balls have little bits of structure. They have electrons going in clouds near the nucleus and they somehow bind themselves together. These things are pretty well understood at some very fundamental level.

Propagation of stress waves, sound, shear waves, whatever you like, gets to be a little more complex than a gas but a little easier to understand in some ways. The complexity comes from the fact that things in different directions look different in solids, and that is going to be the theme of all of this work.

The other part of the complexity comes when the sound wavelengths start to approach the wavelength or the spacing between atoms. At that point lots of interesting things happen to the propagation of sound. The description of the velocity of the sound wave versus its wavelength, the dispersion curve, then becomes a fundamental tool for understanding a lot about the solid.

If you scatter neutrons off the solid and if you do it inelastically, meaning that you throw neutrons that are the result of throwing away most of the energy of one of these giant particle accelerators, you throw neutrons at it, the neutrons have a deBroglie wavelength that becomes reasonable for studying some of the solid-state physics that goes on.

They lose energy in the solid by going in, creating a phonon, which is a quantized sound wave, and off they go, again. If you look at the direction and the energy with which they come off, you can map out the entire dispersion curve from what is called the Brillouin zone center all the

way out to wavelengths of sound that are comparable to the spacing between atoms. That dispersion curve tells you a lot about the binding between atoms and how they are arranged.

Neutron scattering is one of the techniques that provides you with this information at very short wavelengths. So this is the result of a particle accelerator applied to acoustics. The unfortunate thing comes in right here, when you see these error bars. The other unfortunate thing comes in when you realize that in order to really understand a solid I need to know exactly which acoustic mode is doing what.

The neutron-scattering experiment gives me sound velocities to several percent -- that is the slope of this dispersion curve at the Brillouin zone center. But as soon as things get really interesting, that is, where you really want to know what is going on, those error bars take over and you do not know which of these zone-center modes collapsed to zero frequency at the edge.

This is an important point. This mode -- right here (and, of course, you do not know which one it is here, but you can see it here) -- as you change the temperature, the very short wavelength sound frequencies are going to zero. That is the signature of a structural phase transition and these data here for two years messed up the interpretation of the structural phase transition in $\text{La}_2\text{-CuO}_4$, a high-temperature superconductor and one of the first things we decided to study.

The reason they messed them up is the error bars forced everyone to lose track of which mode is which as this phase transition occurred. We are going to fix that -- I think. In order to fix that, we have got to get somebody to pay us to fix this, meaning we have got to get some funding. This is an important part of physics.

(Transparency)

This silly project generated lots of publicity. Has anybody ever applied for an R&D 100 Award? This is an award that a magazine, R&D magazine, has somehow convinced everyone is an important award. It is not, of course, but the magazine makes huge amounts of money from the applicants. Nevertheless, people still apply for these things. It is absolutely true; it is an enormous cost per person.

Our laboratory has a team that sets you up and applies for these awards and we won a few of them. Well, that helps us go get funding and it keeps the project going. You really have to pay attention to these things these days.

(Transparency)

Overall, resonant ultrasound has really done a bunch of interesting things. This is a typical set of things you put down when you are talking to funding people, but we have certainly learned a lot about solid-state physics from resonant ultrasound.

We have learned fantastic things about the electronic states of various interesting forms of matter and we have done it better than the synchrotron radiation guys who do photoemission, which they claim is a microscopic probe. We have used resonant ultrasound, a crude very simple

bit of physics to extract detailed microscopic information about the electronic density of states in solids.

We have also made lots of nice measurements on many different materials. We just published elastic constants for the theorists to puzzle over; it provides a nice data base and it is very easy to do. There are a few other things, but the one at the bottom is kind of interesting, because in the process of getting this technology attached to private industry it has actually become useful and it is being used, and I will tell you about that later, to find cracks in oxygen sensors on General Motors cars.

(Transparency)

I want to begin by talking about elastic waves before we get to resonances and then I will talk a little bit about experimental techniques and what we can learn about them. The key difference here is that we are going to be talking about system solids, where the elastic properties are different in different directions in space.

All of the little laws that we use to generate the wave equation in a gas are still perfectly applicable. We just have to keep track of what direction we are going, and we are going to do that by using vectors and subscripts, and things of that sort. We will also make a distinction and talk about things in terms of strain rather than volume and stress rather than pressure.

The reason is that pressure is considered an isotropic sort of thing, whereas stress and strain at least have a direction associated with them, so we can get a lot of the directional dependence that we are going to need into it just by dealing with stress and strain.

We have a strain tensor, where U is going to be the location of a particle. The variation of that particle position with position in the object is going to be a strain. we want to be a little careful, and the reason we want to be careful here is to keep ourselves from getting bogged down with the math later.

(Transparency)

Let me talk about shear strain for a minute. I want to shear something, so I have got an object and I go "shpptt." Let us say I start with a square object and I push on it, as these arrows indicate. It turns into a parallelogram. I could have pushed on it this way, and this way, and turned it into a parallelogram, or I could have pushed on it along the two corners this way and pulled on it on the corner that way and turned it into the same parallelogram.

The physics at the end is the same for all of these distortions, so it would be very nice to write strains in such a way that a strain that I write that produces this distortion does not depend on the coordinate system. We use this thing called a symmetrized strain.

We pick on the strains so that I add a shear strain this way, and a shear strain this way, together -- that is kind of a crude way of describing it -- to get this very nice symmetrized strain.

This one now becomes somewhat coordinate-independent and it is a measure of the physics of the system.

Obviously, once I have strains, somewhere I am going to try to relate them to the energy of the object. If I stretch a spring, the energy changes. That is really what I want to know. Once I have that stuff written down, I can crank away and get almost any physical quantity I want. I am going to wind up with something where the work, or the free energy, looks like some strain squared.

You can see where these come from. I am going to have some constants and strains squared, and some other constants and shear strains squared to get a free energy. I can separate these things out, because some strains are going to be shear-like and some are going to be compression-like. I can separate this thing out in many, many ways.

One way of separating it out is to pick all the strains that produce volume changes and all those that do not. This has been discussed at length in Landau and Lifschitz better than I can do it here, but I can get the total work on the object as volume-like pieces and shear-like pieces.

It turns out that that is a very important thing to do for solids, because if I know the response to shear and I know one other dimensionless parameter, called Poisson's ratio, which we will discuss later, most of the resonances of lots of solid objects require only that shear modulus and that one dimensionless constant to describe, plus a little bit of geometry.

It separates resonances out and we will use that to great effect later in understanding solids. So this is a useful thing to do and, of course, the thing I really need is a relationship between stress and strain, so I have a stress tensor now, which is basically the response to -- it is a force applied to some surface.

The work is really going to be force times distance (you can work out all these indices later), so that I have something like this. The real complications in a solid are because the solid preserves, or attempts to preserve, some semblance of shape as it strains. If we have, for example, a rod and we pull on it, it gets thinner in the middle.

This is not a really good picture, because if I make a static slow pull on the rod, the entire diameter shrinks. Young's modulus is the relationship between stress and strain when I pull on the ends, and Poisson's ratio is basically a measure of how much it shrinks in diameter as I pull on the ends and make it longer. Conversely, if I were to push on the ends, it would get fatter, except for the corks of wine bottles, which we discussed last night (they have a negative Poisson ratio).

I do not really want to go much further with that, because it is so well covered, and I can never follow these things if I am in the audience in any detail, so we will not do any more than just touch on some of this.

(Transparency)

However, I am now going to give you a whole sequence of viewgraphs on why. Why do we want to study ultrasound? I do not know why you guys want to study ultrasound, but the reason I want to study it is that I am interested in -- actually, I am mostly interested in phase transitions in solids.

One of the wonderful things about specific heat measurements, thermal expansion measurements, and sound velocity measurements is that they show discontinuities at the weakest of phase transitions, a second-order phase transition. They are fully thermodynamic quantities, because they show these discontinuities because of effects associated with the free energy.

The way it works, just to sketch it out briefly, is that I can define the Gibbs free energy across a phase boundary as ΔG . I can take the derivative of that with respect to pressure, and of course I am going to get something that looks like volume -- I will use pressure and volume for the moment instead of stress and strain, so I do not have to have a lot of indices.

If it is a second-order phase transition, the weakest of phase transitions, nothing happens to the volume. So the volume change is zero. But if I take another derivative, if I take the second derivative of the energy with respect to something like pressure, I get elastic constants. Those turn out to be discontinuous.

If you have a jaded experimentalist's taste, you would much rather measure discontinuity than a 1-percent change in the resistivity, so this is the sort of thing that makes you happy, but it is also easier to measure a discontinuity in an elastic constant than it is to make an accurate 1-percent measurement of something else. Part of this has to do with systematic errors and part of it has to do with the ability to pick out the right spot and temperature to see something.

(Transparency)

The way it really works is this. We will do this for a superconductor for the moment. A superconductor has the quality that when it becomes superconducting it throws the magnetic field out. That is different from a perfect conductor, that is, a conductor with zero resistance.

This Meissner effect is tied up with the free energy and, in fact, the very simple discussion of superconductivity says that the free energy difference across the superconducting normal phase boundary is basically the critical magnetic field required to destroy superconductivity squared. I put a volume in here and an 8π in there to match up with various sets of units.

At T_c , at the transition temperature, that is, the transition temperature in zero magnetic field, the critical magnetic field is zero. It is right on the verge of becoming superconducting; it takes no magnetic field to mess that up. It takes the tiniest, tiniest field to mess that up.

This is a second-order phase transition, so I can start to take derivatives of the free energy. I can take derivatives with respect to temperature, so I can get the entropy by taking the first temperature derivative, and you see I get a piece that looks like H_c , the critical field, the H_c at temperature.

I can take the second derivative of that and I get specific heat. As soon as I take the second derivative, I get two pieces in it. This is just simple math. I get two terms when I take the derivative of H_c d H_c dt. One of them is H_c d² H_c dt². The other is d H_c dt. Critical field at the transition temperature is zero, so everywhere on this whole viewgraph every time I see an H_c without a partial or something in front of it, it is going to be zero. So this is zero.

But this is a positive definite quantity, it is something squared. ΔC_p is going to be non-zero. I can do the same thing, instead of taking the derivative with respect to temperature, I do it with respect to pressure. I am going to use pressure here, even though I should have used strain, but let us just do it.

The first derivative gives me the volume change. It has got an H_c here, so the volume change is zero. The second derivative of the free energy is going to give me elastic constants; in fact, in this case it will be the bulk modulus. I find out that the bulk modulus has this d H_c dp term in it, again, squared, so it is positive definite.

I find out that the discontinuity in the bulk modulus (and later we will translate that into particular moduli and show you how you really get lots more information out of it), I have a nice discontinuity in a very measurable physical quantity. So I am very happy studying phase transitions with sound -- even better than specific heat, because the sound velocity is a tensor. I have tons of information in it. Specific heat is a scalar.

If you are going to do thermodynamics on phase transitions, the measurement of the sound speeds is the way to go. I can also take a cross-derivative, by the way. I can take the temperature and then the pressure derivative and I come up with a thermal expansion coefficient, and it does not have any positive definite term in it. It has got two pieces here, but we are already pretty sure that this and this are not zero, so the thermal expansion coefficient also has a discontinuity -- that is a vector-like quantity -- and it is also studied.

However, it is maybe a little more difficult to study, because the changes are smaller, and when you are measuring thermal expansion coefficients, you typically do not measure the expansion coefficient; you measure the length of the sample. So I have a whole bunch of noisy data coming along and maybe it changes the slope, and I have to find the break in slope to get the discontinuity in the thermal expansion coefficient.

That turns out to be a really nasty problem, because your first inclination is to smooth the noise in your data. When you smooth the noise in your data, you smooth the discontinuity in the slope. There are hundreds of papers describing how to massage thermal expansion data to get the discontinuities out of it.

Anyway, you can get kind of an idea of why this is an attractive thing to do for physics. Of course, what is wrong is that I used pressure and volume here and not stress and strain. I ignored

volume and structural changes. If the whole structure changes at a phase transition, you have to handle that carefully.

I have also ignored magnetic-field effects, which do funny things, and I will show you how that works in a minute.

(Transparency)

So what we want to do is talk about elastic constants and connect them up to fundamental physics as best we can. This is kind of a true statement, and the only reason it is true is because many people do not stare at elastic constant matrices for a long time and realize how utterly simple they are.

It is like when you get your Ph.D. Everybody thinks they take you into a little room and they suddenly tell you all these secrets and you are never allowed to repeat them to a normal person, again.

(Laughter)

The same thing goes for elastic constants. There is really not much to them, and we will see why in a minute.

The reason things get complicated, as I mentioned before, is that when I push on things this way, they bulge this way. If the insides of the solids have complex molecular structures, not only when I push on them this way do they bulge this way, they might bulge that way in some very weird direction. All we have to do is keep our wits about us and handle these things properly and we can deal with them quite nicely.

Of course, a gas does just the same thing. If I push on it this way, it tries to bulge that way. I am going to make a comment -- you will see what orthorhombic means in a little bit -- that an orthorhombic solid is one in which, when I push on it this way, the only stresses and strains that result are either parallel or perpendicular to the applied stress.

In that case, I do not get a lot of cross-derivatives or whatever you like, and I can get a nice set of strains from this thing. I can compute things like Poisson's ratio if I know what these strains are. What is Poisson's ratio for a liquid? It is a half, right, because the volume stays fixed, approximately (unless you are using it in a heat engine or something of that sort).

In a liquid, when I squeeze it this way, if I put it in a loose plastic bag, when I squeeze it this way, change its length this way, it bulges this way, such that the volume stays fixed. That gives me a Poisson ratio of a half. Solids do not have Poisson ratios of a half; they have Poisson ratios that are less than a half, because the solid itself attempts to prevent some of this bulging.

(Transparency)

We want to study elastic constants because they connect up in very complicated ways to details of the solid. The more complicated the way they connect up, the more we are going to

learn. We are really going to be studying the free energy, and there is an entropy piece to it, and then there is a simple stress-strain relationship.

The stress-strain relationship is Hooke's law. The way it works is that I am just going to keep track now of all the directions and space that I have. I can provide stresses and strains. I have already made a mess of this. I was thinking of showing you the whole ball of wax here, but I decided against it.

Look, remember that I have a solid and I can apply a compressional stress, that is, a normal stress, this way, this way, and this way -- that is three. Then I can apply a shear stress this way, this way, and that way -- that is another three. But there is another corresponding pair of shear stresses that go -- for example, if I applied a force this way on this surface, I get one shear stress. If I apply a force this way on this surface, I get a different one. Physically they are the same.

But when you start to write this whole thing down, you begin by treating those as different ones, so I have three compressional stresses -- boom, boom, boom -- and six shears, which gives me nine things that I have to deal with. That means that, in reality, if you want to write down a matrix that transforms properly, meaning I can do simple geometry to find out how the stresses and strains look in a different direction in space, the thing is really a 9×9 .

However, the simple physical symmetries that you would like to apply make it symmetric around the body diagonal of that matrix, and that reduces by a lot the number of elastic constants. Remember, the elastic constants are going to be the entries in this matrix that connect stress to strain.

Then, if I use this symmetrized stress, where I say that this shear and this shear do the same thing physically, I get rid of some more. In the end, I am down to 21 different entries. Twenty-one different entries occur when I have a 6×6 . So I have the six different ones along the diagonal and then all the guys off this diagonal are the same as all of these and, if you add it up, it is 21. That is where you get the 21 maximum possible number of elastic constants.

That gets hard to see and hard to do anything with, if we want to teach it a little. When we are going to use this matrix to find out sound propagation, the way the solid compresses, and that sort of thing, I am going to do things like take this times this, plus this times that, plus that time that, plus that times that, and so on, right? That is how you multiply a matrix times a column vector.

These three I am going to call shears. For example, e_4 is really $dU_1 dX_2 + dU_2 dX_1$ divided by 2. It is half of this shear plus half of this shear.

If I have an orthorhombic solid, that is the solid in which everything happens either perpendicularly or parallel to the applied stress, it turns out all these entries are zero. It becomes very easy to understand, because when I squeeze it this way, I do not have things shooting off at

some weird angle. They are going to either bulge or they are going to shrink, and when they bulge or shrink I need only these entries.

If this were a highly symmetric object elastically, for example, a glass, it turns out you need only two elastic constants, and so there are specific relationships between these. But you can see very easily that a shear stress with an index 4 here connects only -- the only time I am going to see it is when c_{44} multiplies e_4 .

In an orthorhombic solid the shear stresses and strains separate out beautifully and they look just like Hooke's law for a simple spring and you are done. The other ones, the volume ones, are a little more complicated, because they tie up with Poisson's ratio, and I will show you how.

(Transparency)

The question we want to ask is how do you use this mess? Let us compute the bulk modulus for an orthorhombic solid. Bulk modulus is going to be the response of this thing to hydrostatic pressure, so I make the stresses here equal to something.

What I am doing is I am saying -- remember, this entry is the x or the one-direction stress. This is the 2, or the y direction, this is the 3, or the z direction. If I put the thing in water and I compress the water, the stresses are the same in all three directions, and there are no shear stresses. So I put pressure, pressure, pressure, zero, zero, zero. These are the shears.

I know right away that if these are zero, these are going to be zero, because this is a little diagonal piece and there are zeros everywhere. Not surprisingly, when I apply hydrostatic pressure, I do not get any shear stress or strain. I am left with a little 3×3 here multiplying these three, so the problem is reduced now.

Instead of the original, full-blown 9×9 matrix, all I need to do is deal with the little 3×3 , and that you can do with your pocket calculator, if you want to. What I want to do is I want to find the volume change when I apply hydrostatic pressure. Absolutely simple.

I know that the strains to first order, the volume will change by roughly the sum of the three strains. These are dimensionless, by the way, strain is dimensionless, it is $\Delta L/L$. I want to compute $e_1 + e_2 + e_3$, and I just do it. For an isotropic body, it turns out that the two constants that Landau and Lifschitz like to use are the Lamé coefficients, one of which is the shear modulus μ , and that is just what you think it is.

The other one is a thing that they call λ . The way it connects to science is that λ plus 2μ is equal to c_{11} . What I have here is λ plus 2μ and I have λ , λ , so on. λ plus 2μ defines the longitudinal, or compressional, wave sound velocity speed in a solid, and μ defines the shear wave speed.

For the purposes of this discussion I am going to substitute these into here, do this simple multiplication out, solve the three sets of equations, which have a lot of symmetry built into them, and I wind up with a bulk modulus, which is $V dp/dV$ is λ plus $2/3 \mu$.

It looks like, because this μ piece comes in, that there are shear effects in it, but there are not. It just happens to be an artifact of the way these elastic constants are defined.

As we get to a more complicated solid I am not going to be able to do that so easily, because I will not have all those zeros.

(Transparency)

The general prescription, by the way, is I have these strains that look like this. This might be the x, y, or z axis and this would be x, y, or z. I have these pieces and each of these little strains has essentially an index i and an index k.

The way I got that 6 x 6 from the 9 x 9 was to collapse them like this. Whenever I have a pair of indices, 1-1, I call it 1, 2-2 is 2, et cetera, 2-1 and 1-2 are 4, and those are obviously symmetric, because they relate to shears. So I now, in the general solid, have a mess, and it is really a nasty mess.

I have stresses related to strains via Hooke's law, but I have an entry in every one of these locations. Things get even worse, however, if you apply a magnetic field, when this symmetry between du_i/dx_k and du_k/dx_i disappears. In a magnetic field, especially a high magnetic field, we can even lose that symmetry that reduces it to 21 and get more elastic constants.

This is the generic problem that if we have bad luck, meaning the physics system we are interested in is less symmetric than orthorhombic -- orthorhombic is the one in which all of these are zeros, the simple one to see. Orthorhombic symmetry is the general case of a thing where it is intuitively simple to see what is going on. Cubic is a special case of orthorhombic. Isotropic is a special case of orthorhombic.

But if we apply a magnetic field, every one of these can be different and that can be really nasty.

(Transparency)

Just to let you see for a second where some of the differences come in, if I had a solid in which each of these little dumbbells is located on a cubic lattice, and I squish it like this, those little dumbbells will rotate as well as compress. You can see that I can get strange strains for a simple stress.

(Transparency)

The way the whole business goes is as follows. Remember that orthorhombic, I will say it, again, when you squeeze it like this, everything happens parallel or perpendicular. However, when I squeeze it this way, things are different from when I squeeze it this way or this way.

If I increase the symmetry to what is called tetragonal, tetragonal is just like orthorhombic, except when I squeeze it this way it does one thing and when I squeeze it this way or this way, it does something different, but they are the same. It is a lattice that has a square base and a longer or shorter vertical direction.

When I go to tetragonal, instead of needing the nine different numbers I need for orthorhombic, if I increase the symmetry to tetragonal, I need only six moduli.

The next one is hexagonal, which is really a wonderfully interesting one. Hexagonal is one in which the tetragonal lattice has a particular constraint. If you work out the fundamental properties of a hexagonally symmetric lattice, the lattice is really a bunch of prisms that look like hexagons on the base and they stand up.

It turns out that the sound speeds are isotropic in a hexagonal plane. They are not different this way from along a diagonal; it is absolutely isotropic. That is very important, because many physical systems are somehow constrained to be uniaxial. For example, if you stretch out and put a whole bunch of fibers into a plastic matrix, it is different this way from this way, but it is the same all around. That has hexagonal symmetry and you need five elastic constants to describe that.

The next one down from that is cubic. In cubic it is the same this way as this way as this way. There I need only three elastic constants. The interesting thing on the cubic system is that I have perfectly good shear waves propagating parallel to one of the cubic edges and also along the diagonal. So in the cubic system I start to make contact with the simple solid-state physics course that says you need one compressional and two shear waves to describe a solid.

In the cubic there are two clean shear waves that travel at different speeds. One is along a diagonal of a face and the other is along the edge. And I need only three elastic constants.

Finally, I get to isotropic, in which I need only two. There is a shear wave speed and a compressional wave speed and that is it. But from orthorhombic on down, I have zeros everywhere here and it makes a lot of computations very easy to do.

DR. MAYNARD: What about the constants?

DR. MIGLIORI: Are you going to tell us later? Do you think it is two?

DR. MAYNARD: It is two, yes.

DR. MIGLIORI: That is what I think, too. It kind of makes sense, because in a long-range picture it is very hard to see the short-range order, and so it has to get isotropic-like. If Dr. Maynard says it is two, it is two.

(Transparency)

This is to show you that when you finally try to write down a wave equation -- here is mass times acceleration and here is the force. When I finally write down a wave equation, it is a total mess. It has got lots and lots of subscripts all over the place and, worse still, most waves in a solid are neither transverse nor longitudinal.

I get transverse shear waves in a cubic crystal when they are propagating parallel to an edge or along a face diagonal, but a wave, any wave, that propagates in another direction no longer has

the displacements either parallel or perpendicular to their direction of travel. Because of that, all the wave equations get messy and I have to be quite careful about how I deal with them.

DR. SWIFT: Has anybody tried to get useful physics information out of the imaginary part of all those elastic constants by measure the Q's and stuff like that?

DR. MIGLIORI: The answer is yes. I will show you data on the measurement of ultrasonic attenuation using resonances via the Q. But in addition, that elastic tensor has been worked out where one puts a complex element in each of the positions in the matrix and you can actually get not only the Q, but you can tell which elastic constant is responsible for the dissipation and how.

Therefore, you can connect that up later to the fundamental physics, which atom is squiggling around causing the dissipation. That has been done and it has been done by the Japanese mostly.

We have actually found that a lot of physics can be extracted without too much complication. That I will talk about a little bit, that is, how you can get at attenuation and know what is causing it without really solving two 9×9 matrices.

We know that solids are different and we know why.

(Transparency)

We have a rough idea of why elastic constants are interesting things to study if you are interested in solid state physics. I want to talk now about the general quality of resonances as applied to this.

This says why resonant ultrasound? There is this thing that rich people say. They say, "The only people who say that money isn't everything are those people who don't have money." Well, I did not have a pulse-echo system --

(Laughter)

-- and just thinking about trying to get 60K in capital equipment from the DoE 10 years ago to buy all the pieces you would need to put together one like Moises', it just made me tired thinking about it.

The only thing I could do was to measure resonances. It was only much later that I realized that that was not only the cheapest but the best thing to do in many circumstances, that it worked far better than pulse echo. Pulse echo is the technique that requires that you send a tiny burst of sound off into space, space being a solid, bounce it off the edge, and time its arrival. From that you get the sound velocity.

You do not even need a pocket calculator. You can do the arithmetic in your head to get the speed of sound.

However, there are lots and lots of subtleties to this. Remember the bar? When we stretched the bar we got Young's modulus. We did not get the longitudinal sound wave speed. So you

have to be careful when you do pulse echo. You have to be careful that the wavelength is not too big, or you are going to wind up with Young's modulus instead of c_{11} .

You have to be careful that the transducer does not cover the whole face or else lots of strange sound waves propagate down the edges. What is happening here, the reason I do not get Young's modulus sound speed in this sort of pulse-echo measurement is that all the undisturbed material around this sound beam is constraining the strains from that sound wave so that I really get plane waves, for the most part, and not Young's modulus waves in the system.

I also need very fast electronics. Every time I hear "fast electronics," it makes me nearly ill, because fast means wide bandwidth, which means you are opening all your doors and windows to noise. At some level you are going to wind up measuring tons and tons of noise that you did not really need.

However, if you are really confused about a system, there is nothing like pulse echo, because you do not need computations, and you will see how bad the computations are for resonances. But there are lots of subtleties and pulse echo requires fast electronics, wide bandwidth, lots of noise. You have to be very careful that the experimental situation is what you really think it is in order to get good data. However, when it is done well, lots of good things have come of it. I did not have money to buy one. The real reason is we were broke.

There are a couple of interesting things about resonances relating to transmission. Let us say that I am in Monterey, California, and I shine my laser pointer up into what they call air here -- (they do not breathe air they cannot see).

(Laughter)

I am firmly keeping in mind thermodynamics, so I shine the laser beam up in the sky. What does it do? It fuzzes away into the fog, yes? Nondissipative. No energy has been lost, it has been scattered. Scattering, for the solid state physicist, is primarily the result of laziness. You did not make the material clean enough, you left grain boundaries in it, you left junk in it that scatters sound. Very rarely is there any physics tied up in that stuff.

If I want to measure the intrinsic thermodynamic losses associated with changes in structure or changes in physics, I do not want to know about scattering. In the pulse-echo system, this little flashlight beam of sound is going to scatter all the defects and it is going to fuzz out. If I measure the intensity of the echoes as a function of time, I am going to see scattering effects.

I am also going to see effects associated with the fact that maybe these faces are not parallel, so the beam starts doing this, and it walks off the received transducer. I am also going to see diffraction effects, because I have a finite-sized antenna, aperture, lens, whatever you want to call it, and I get diffraction.

If I want to use this technique to measure ultrasonic attenuation, I really have to keep my wits about me. But if I do a resonance measurement, the resonances -- for the purposes of teaching, I

whack this thing and it rings. The ringing, the decay of the ring, depends only on the intrinsic thermodynamic dissipation in the solid.

When I make resonance measurements of the attenuation, I see, really, the dissipation, and I do not see scattering and I do not see beam deflection, diffraction effects, and so on.

The result is that in many systems pulse echo requires the subtraction of several large corrections to get a small answer, whereas resonances produce, really, the losses and nothing else.

Dr. Maynard did this wonderful thing in measuring the attenuation in optical glass for fibers. He took a big chunk of glass and he drove it on one of the resonances of this big chunk with a laser beam. The only thing that caused that thing to ring when he chopped the laser beam were the real losses, which is what they wanted to measure.

That was a system where the losses were so low that there was no hope of finding them by simply doing a transmission experiment. This is glass where they wanted to go for kilometers, so you could not really do it very well without a resonance measurement.

(Transparency)

Furthermore, if you get to do resonances, you get to really clean up the electronic system. We can tailor, in a resonance experiment, lots of stuff, because primarily we are taking data all the time. In the pulse echo system you are hardly taking any data at all. You send off a half-microsecond pulse and most of the time you are sitting around waiting with all your electronics wide open to noise.

This is a somewhat biased signal-to-noise comparison of a resonance measurement versus an impulse. The resonance measurements are really are a swept sine measurement. I can begin with the drive power per unit bandwidth. For an impulse I have extremely high peak power, but the full bandwidth of the system has to be not the frequency of the sine wave that I am transmitting; it has to be determined by the time resolution with which I want to know when that echo comes back.

If I have a 10-MHz sine wave and I generate a burst but I want to know the sound velocity to 1 percent or 1/10 percent, I may need to have a nanosecond response time for my electronics, so the noise bandwidth is enormous. In fact, the power-per-unit bandwidth, even though the peak power is fantastic, the power-per-unit bandwidth for the real measurement bandwidth is typically lower than a swept sine system.

You might put a kilovolt burst on a pulse-echo transducer, the system back there might put a volt on a resonance transducer, and we still wind up with more power-per-unit bandwidth on a drive. The noise bandwidth for the complete measuring system, then, might be 10^9 Hz for a pulse-echo system. For the resonance system, if I take all the modes I know in advance I want to measure and I measure them five times wider, or something like that, than their actual resonance width -- so now I am trying to find every place I can to add noise bandwidth to my resonance

system, because it looks too good otherwise -- even when I do all of that, I wind up with only, say, maybe 10 kHz of real noise bandwidth for a typical resonance measurement.

Then I have duty cycle. This is very generous for pulse echo. Of course, for swept sine I am sitting there all the time measuring. Finally, I take the square root of all of these things and I find out that I am in pretty good shape with a resonance measurement, especially if I am worried about noise or have small samples. So if you can do a resonance measurement, do it. If you can get the physics out, do it.

I have got this rubber baseball bat here, Moises. You can beat me now. Go ahead.

DR. LEVY: These two values, how well can you measure velocities, to how many parts per million? I can do 1/10 of a part per million.

DR. MIGLIORI: Yes, I can do that.

(Laughter)

DR. LEVY: About 10^7 , maybe?

DR. ATCHLEY: Absolute or relative?

DR. LEVY: Relative.

DR. MIGLIORI: We can do roughly 10 times better than pulse echo on absolute. An average resonance measurement is 10 times better than the best pulse-echo measurements on absolute, because we do not have to make transducer-bond corrections and other things like that (you will see that later).

On relative measurements of the sound speed it turns out that they are comparable. In fact, pulse echo, for very strange reasons that I do not completely appreciate, may be an order of magnitude better than resonances. The Bell Labs guys have reported numbers in which they see variations in sound velocity of a part in 10^8 .

DR. LEVY: So do we.

DR. MIGLIORI: We have made measurements where we can see a part-in- 10^7 change. Unfortunately, I do not completely understand the fundamental physics associated with those numbers, because all of the systems are dissipative and, depending on the level of dissipation (I will hit that in a minute), a part-in- 10^8 change in the sound speed can come entirely from changes in the losses.

When you get into that circumstance, and I will say this again in a few minutes, it is not clear what elastic constants mean. If the dissipation is high compared to something that you care about, then you had better not fool yourself into believing that you have sound propagating with attenuation.

What you have is a system in which there is a relationship between stress and strain. It is not Hooke's law and it is heavily tied up with dissipation. What you really need to do is look at the

relationship between stress and strain somehow and not pretend that the sound velocity and attenuation really describe the system well. I will show you that quantitatively in a minute.

DR. LEVY: How about attenuation?

DR. MIGLIORI: That is where it is all tied up. If the attenuation is low, you can do --

DR. LEVY: I mean, what is your sensitivity?

DR. MIGLIORI: Well, yes, our instrumental Q's are about maybe 5×10^5 or so.

DR. LEVY: In the best of cases.

DR. MIGLIORI: Yes, in the best of cases, sure. They are all roughly in the same ballpark. However, I can do it on a sample the size of the head of a pin, Moises.

(Laughter)

If I had Moises' pulse-echo system I would use it all the time as well as the resonances.

DR. HOFER: We were talking a little but about pulse echoes last night when we were here and you were making a comment about bandwidth and resolve time. Suppose you have something like a 10 MHz electronics and you were trying to resolve in a few nanoseconds or something.

Of course, you cannot do that if you just have some arbitrary starting time and you want to measure the time of the length of the first pulse. But if you have multiple echoes, you can sample first one echo and then the next echo. Can't you do lots better in the bandwidth of your electronics by just matching the time delay between the two?

DR. MIGLIORI: But of course. Absolutely. In which case, the more you push that, the more echoes, and the more you push that, the closer you get to a resonance measurement, if you think about it for a while.

The way you get high precision in a pulse echo measurement is you forget about the pulse shape and you look at the carrier -- maybe it is a 500 MHz carrier -- and you bring it in coincidence with gates. You look at the starting pulse or maybe the sixth echo and then you time delay the seventh echo and you bring it in and you use a mixer, average it, multiply it, and make it zero, and then start twiddling the frequency by parts in 10^8 to hold that zero as you change temperature. That is about as close to a resonance measurement as you could get without actually calling it one.

(Laughter)

Do not worry, Moises knows I am jealous, so it is okay.

DR. LEVY: Well, I am jealous of you, too.

DR. MIGLIORI: See, Moises did this right. He did his pulse echo system -- he did not buy one -- he bought every piece of it and put it together, so he actually knows what he is measuring, which is not the case for a lot of commercial pulse-echo systems. You can believe his data.

(Transparency)

DR. MIGLIORI: So we are going to measure resonances for whatever reason, finally, you like the best, and we are going to do it the dumbest possible way, just because I have total control

of the noise by simply driving the system with a synthesizer, amplifying it, mixing it down so I can design a filter that has a fixed center frequency, convert it, computer, there we go.

A huge part of the story is tied up in this block diagram. By the way, here is one of our low-temperature cells with a 2-mm rectangular parallelepiped sample in it. That is typically how we make measurements. We just drop the sample and there is a gram, maybe, of force applied to the points of the corners. Look, we are trying to do resonances, right? I want to get all the resonances, which means I want the excitation of the object to have as low a symmetry as possible, so I excite all the modes. So you hit the points of the corners.

By hitting the points of the corners, number one, I have very weak coupling to the transducer, so I am going to have to do the electronics design perfectly. We are going to wind up with thermal noise-limited electronics. In exchange for that, I do not need to put any goo on the sample to attach the transducers. It is dry-point contact.

That effect means that the only thing I have to include in the computations, if I do not squish this too hard with transducers, is the sample, with no corrections. There are no corrections. In determining elastic constants with this technique, there is this business of how do you know when a cesium clock is accurate or, as in the ancient Chinese proverb, "Man who has one watch always knows the time. Man who has two watches never knows the time." But he knows how badly he is off.

In fact, this technique has so few corrections that we have come to believe that the only test of its accuracy is to make several measurements on several samples cut differently and look at those as the error bars, because all of our corrections on all of the other systematic errors appear to be less than we can resolve. They are just not there.

DR. MARSTON: If those corners are sharp at atomic dimensions, does the static stress diverge at those corners?

DR. MIGLIORI: Yes, sure. What happens is we do the best we can, you will see in a little bit, to make those corners really sharp. You put the transducers on them. They yield. They yield until the surface area of contact equals the one gram per unit of force we apply to it.

At that point everything stops. But that is a pretty sharp contact, it turns out to be. Therefore, the corners are infinitely soft in any reasonable limit, meaning their mechanical impedance is very, very low, because it takes very tiny forces to change the distance over which those transducers are separated.

Because those corners have mechanically low impedance, or are soft, it means that to about a part in 10^6 or so we have an object with free surface boundary conditions. That is critical to the analysis that we will do later. This is an object suspended freely in space to about a part in 10^5 or 10^6 .

There are lots of things going on here. One of the things I want to do is measure small samples. Small samples have high resonance frequencies. I do not want to measure the resonances of the transducers, so there is going to be some art in the transducers. I am going to tell you about what we did, and you have read about Dr. Maynard's approach, and others.

The transducers are small. The point contact means the signals are small. We are going to have be very, very careful about the preamplifier, and I am not going to tell you about that, unless you ask me later. Once we are out of the amplifier and preamplifier, we are kind of home free, except that, look, you know, you have a resonator, if it has a Q , the Q tells you how long it takes you to ring up and ring down, again. That absolutely defines how fast you can make a resonance measurement.

I am going to wind up taking, for example, on a typical run from room temperature down to 4°K on a sample I might be interested in, 40 or 50 megabytes of data. I have to do this as fast as Mother Nature is going to allow, so we need to be able to have every part of the electronics go as fast as it is reasonable to go.

In order to do that, we designed a dual digital synthesizer. One output channel provides a local oscillator that is 1 kHz separated from the drive frequency. We have 32-bit control over relative phase and frequency and 16-bit control over amplitude, and you can put that in an assembly language program and make that synthesizer step perfectly quickly.

Furthermore, by using a true digital synthesizer, we do not have phase noise and other garbage to come out of things like these very fine Hewlett-Packard frequency synthesizers. They just do not work as well, because they use phase-locked loops. These are all details you can ask me later; I do not want to go into a lot of it.

One of the things that we had that is a disaster in most systems like this is the filter that comes from the mixer. The mixer multiplies the local oscillator and the drive. The difference frequency in our case is 1 kHz. I want to measure the amplitude of that 1 kHz and put it into the computer. Now, how do I do it? I rectify it, put an R.C. filter on it, measure that with a DC voltmeter -- zap! You are dead.

There are very clever digital signal-processing techniques that come in right here, where I have to deal only with a 1-kHz signal instead of a 2-MHz signal that allows me to build an R.C. filter that settles in exactly 1 R.C. time constant, always, with no additional noise. We have used that, so that we actually can sweep through these resonances at rates that are determined by the Q , not by the electronics.

STUDENT: What is the number of steps like this?

DR. MIGLIORI: Here? Okay. I want to drive resonances. Resonances have symmetries. For example, some of them might go like this, some like that. If I put transducers on the faces of the sample, number one, I load the sample. It is not a free object any more. Number two, I will

accidentally drive it at high-symmetry locations, which means some of the modes will not be excited.

So by hitting the points of the corners, it turns out that only under the very weirdest circumstances are there nodes there; that is, you have to have a bunch of accidents occurring with elastic constants and relative sample dimensions to get a node there on only one mode.

DR. BASS: What is the typical time between steps and frequency?

DR. MIGLIORI: Take FeSi. A 2-mm cube of this stuff has a Q of 10^4 and a resonance frequency of a megahertz. So it is 100 μsec to ring up that resonance. That means I had better take on the order of 100 μsec to make this measurement. It turns out you need two or three times that.

But I can do it either by knowing its resonance frequency, going right to that and waiting 200 μsec to get the amplitude, or I can take 200 or 300 μsec to step through the whole thing very, very much more rapidly than that. It works out to be the same.

We could go through that whole resonance in a millisecond, and typically we will go through it in 10. We will slow things down. Our electronics is fast enough to step it about 100 $\mu\text{sec}/\text{step}$, if we choose to.

DR. MARSTON: Presumably, though, if you want to measure the Q itself and you interested in dissipation, you have to slow down that process, is that true?

DR. MIGLIORI: We have to slow it down enough that we get the right answer. And you know what the right answer is going to be, because you can do it in successive approximations.

(Transparency)

This is what we are dealing with. I mean, does this not warm the heart of a physicist?

(Laughter)

Thirty Hz. Lorentzian resonance shape with a little bit of coherent background. This is your answer, Moises. How well can I tell the center frequency of a good low-noise Lorentzian peak like this? Certainly 1/100 of the line width. Maybe better.

So the result of a lot of care gets us resonances like this quite often, lots of them.

(Transparency)

I want to talk about resonances for just a minute before we proceed. Here is a simple mass spring resonator. I have got two masses in this because I do not want the proverbial lab frame to screw things up. I want this to be a free object that is resonating. And I have a damper, a little piston with some leak in here. I have M , K , and some damping constant.

(Transparency)

We do the usual things to this and out comes an equation of motion and I solve it by an harmonic approximation and I get a complex frequency. The complex frequency means that when I multiply this by i , one part of it is real, which means it decays in time.

The decay time is something tau. I can write it in lots of ways. Here is one way of writing it. I can solve this problem. Free resonator, you pull it, you let it go, it goes boing, boing, boing.

The Q is $\omega \tau$ over 2. What is the Q ? It is about the number of cycles this thing rings for in very crude circumstances. More accurately, it relates to the decay constant, a dimensionless decay constant, for an exponentially decaying ringing oscillator.

I can also drive it. I can put a force on one of the masses, connect it in some physically reasonable way, and drive the resonator and sweep it through resonances. When I do that, I get lots of things happening.

(Transparency)

I get the amplitude-of-motion, without regard to phase, peaks at the resonant frequency. There is a real part and an imaginary part that also describe the motion and these basically describe the phase shifts between the driver and the response. The phase shifts by 180 degrees as you go through resonance.

(Transparency)

So what is the resonant frequency? The displacement is a maximum when ω is equal to ω_0 , which is the square root of K/M times $1 - 1/(2Q)^2$. That is not quite the frequency of oscillation that you got when you took it, let it go, and allowed it to die away. There is a piece in there, there is a real and an imaginary piece to the frequency of that, and the real piece is $1 - 1/(4Q^2)$.

It turns out that the phase is exactly in phase with the driver when ω equals ω_0 , when the thing looks purely dissipative. So what is the resonant frequency of this thing? I do not know. And if Q is changing, you do not know, either.

So when we have a dissipative system, we have to take some care to understand the effects of the dissipation and the details of how we make the measurement and what we chose for the resonant frequency. Very often, in a solid state physics system, the elastic constants do not change a lot, but the dissipation goes nuts.

I can show you at least a few publications on high-temperature superconductors in the last eight years or so where people have made measurements on polycrystalline samples -- not Moises -- and they have mistaken changes in dissipation associated with intergranular effects for physics, a change in the sound velocity, because the Q was not 10^5 , it was 100. Then you have trouble. You can see that I can easily get into trouble when the Q starts to hit 100 or lower.

An analysis of that little mass spring-damped oscillator is also almost everything you need to know to design the bandwidth and the electronics. That little mass spring oscillator tells you something about how well frequency is defined and, also, what bandwidths you need to use to make a resonance measurement. You can see it from the graph.

That tells you exactly how to design the filters and the R.C. time constants. The filters, a nice single-pole filter, acts just like a damped mass-spring resonator. You know what the filter response is and how long you have to leave the signal on the filter by analyzing the little mass spring resonator. You can use that to back yourself up to noise bandwidths and so on.

This is almost a global problem, so not only does that mass-spring oscillator tell you how fast you can measure; it tells you how to design all the pieces of the electronics to make the measurement allow just the signal you want and none of the noise.

(Transparency)

A little more on resonances. If you have a solid object, the resonances are the eigenfrequencies. What does that mean? It means it is a set of frequencies at which wave equations do nice things with their phases at the ends of the objects. If it were to look at a system like this, where I have one material here and one here with pulse echo, I get a nice reflection off of this point. I know where this is. I get another reflection this way.

However, if I look at the resonances of this, it turns out that this object is not a lot different from a homogeneous object in terms of resonances. The resonances do not see this boundary very well; you just get new resonance frequencies for the system with different boundary conditions.

You can extrapolate this a little bit to note that if I take an object, like a rectangular parallelepiped (that is a cube with all of the edges different -- not all of the edges, some of the edges), that object has a set of mechanical resonances. If I made it of an isotropic material like glass, I would have a set of resonances, which has two elastic constants. Now it is a rectangular parallelepiped, not a cube.

Now I go from glass to some orthorhombic solid with nine different elastic constants. I do not change the number of resonances. None of them were degenerate before, none of them are degenerate afterward. The actual positions of the resonances will be different. Those tell me about the elastic properties of the solid and the model I used to describe it. That goes for this sort of thing, too.

(Transparency)

The solutions to this problem, as you all know, come by the pulse-echo technique. I bounce a whole bunch of little waves back and forth that match up a whole bunch of boundary conditions and wind up with, in this case, an implicit equation for the K vectors in the two materials. I note that ω is constant everywhere. The frequency, and that is an implicit equation for frequency, so if I know densities, sound velocities, and lengths, I can compute the resonance. Conversely, you can see if I know the resonance and I measure the lengths, I can get these sorts of things.

I just get new resonances, that is all. You should remember that throughout, that you can do all sorts of strange things to a resonator and still have perfectly good resonances. It is the math that

you do afterwards that tells you what it is that you measured, unlike pulse echo, where you maybe just need your pocket calculator.

(Transparency)

I am still backing into this slowly here. How well do mathematical models and resonance experiments with well-designed transducers and electronics work? The most precise commonly made physical objects are ball bearings. Routinely these are produced with on the order of 10 ppm mechanical shape errors.

This is a resonance measurement on a silicon-nitride ball bearing. They are used in fuel pumps for space shuttles and other weird things. These ball bearings cost \$300.00 each. They are polycrystalline and very round, so we can make measurements of their resonances. These are the computed frequencies based on something I will show you in a little bit, a very good computation. These are percent -- let me say it, again -- percent errors between measured and computed resonances.

This is the highest absolute accuracy measurement of the elastic constants of polycrystalline silicon nitride ever made, and took about half an hour. It is so good that if you breathe on the ball bearing you can see all the resonances shift from a 10th-of-a-degree change in temperature. You can see 1 mK at room temperature. Look at that! Isn't that something? We should just sit here and admire it for a little while.

(Laughter)

DR. HOFER: You made a big deal of the sharp corners of your cubes and mounting it. How do you mount a sphere?

DR. MIGLIORI: You are just really careless about how you align the transducers and let them touch.

DR. HOFER: No sharper points anywhere?

DR. MIGLIORI: No, it is still point of contact and it is very light loading.

These were some of our tests on loading effects. We took those data and started loading the ball and until we got up to about 10 gm we could not see any shifts in anything. Basically what we have done to determine the loading is to use the mathematical models, fit the data on very well-defined shapes and reduced the loading until we hit an asymptote and it is done.

DR. GARRETT: On the legend for table No. 1, you have this μ , which is 1.2 something times 10^{12} and the sigma has no dimension.

DR. MIGLIORI: Right, that is Poisson's ratio. On a ball bearing every single mode is described by the shear modulus and one dimensionless parameter, the Poisson ratio, and the whole problem scales for all ball bearings. That is why shear modulus is really -- shear modulus and Poisson's ratio are the nicest things to use in elasticity.

DR. GARRETT: I was confused by the fact that I had never seen a Poisson's ratio to four digits.

DR. MIGLIORI: Eat your heart out.

(Laughter)

(Transparency)

More routinely, we can do simple isotropic things that are kind of messy. Yes, go ahead?

STUDENT: Is it possible to measure the.... frequency instead of continuous frequency?

DR. MIGLIORI: I am looking at resonances, so I have to measure many of them. The resonances are defined by the object. If I wish to do the frequency dependence of an elastic constant, I simply look at several objects parking the resonances and the frequency ranges I wish to study.

DR. HARGROVE: I think maybe the point is that if you determine the losses from the Q, that is a loss at a particular frequency.

DR. MIGLIORI: But of course.

MR. HALLEJ: You need to know the density and the Poisson's ratio before you do these measurements?

DR. MIGLIORI: Just the density. We measure Poisson's ratio.

MR. HALLEJ: So your density measurement has to be as good, then, as your other results?

DR. MIGLIORI: Right, if I have an object whose dimensions are defined to 5 ppm, that is how well I know the density.

MR. HALLEJ: Okay, because in the pulse-echo method that is one less measurement you have to make, right? You do not need to know the density to measure the speed of sound.

DR. MIGLIORI: Right, but I want elastic constants, not speed of sound.

DR. HARGROVE: Elastic constant is always a modulus over a density. You still need the density.

DR. MIGLIORI: Yes, there is no free lunch. I can get the speed of sound immediately without the density by measuring the dimensions of the ball bearing, which I need to know anyway. That is equivalent to the pulse-echo measurement. If I want elastic constants, I have to know the density and so does Moises.

MR. LAVRENTYEV: If you are trying to find the elastic constants from the resonance positions, how precise should the shape be? If your ball is a little bit off --

DR. MIGLIORI: The errors are linearly proportional to dimension square root of the elastic constants. If you want a part per thousand absolute accuracy, you had better have part per thousand accuracy in the dimensions.

It turns out we can do a little better than that, because we can fit two of the three dimensions, which means we get a square root of three on the error bar, on the real random error bar, because I can measure three dimensions and I can fit their ratios as well as all the elastic constants.

It turns out that the errors correlate, so it is as if you knew that one of the dimensions about the square root three times better than your measuring instrument tells you. It is a minor effect at this point. At this point the data are good enough for almost anybody.

MR. LAVRENTYEV: The geometrical size.... You know that it is close to a sphere but you can --

DR. MIGLIORI: I will show you how that is worked. We have done a perturbation analysis of a sphere in some detail.

(Transparency)

This is more typical. This is a chunk of steel from a Saturn drive gear in the transmission. It is fine-grained polycrystalline steel. It is moderately dull -- moderately dull? it is a cure for insomnia if you are a solid state physicist -- but it pays for lots of high-temperature equipment.

(Laughter)

This is typical of what we will get for RMS errors in the fitted and calculated frequencies. However, I will show you in a minute that that is not the error bar in the elastic constants. What we are doing, remember, is we are measuring resonances. We are fitting this. We are fitting a few elastic constants to 20 resonances and maybe the dimensions.

When I am done, I have an N dimensional, where N is the number of elastic constants surface, and that surface I am looking for a minimum in the error function, that surface is the error function, and it will have principal directions of curvature. In each principal direction the curvature will be different. The curvature tells me the error bar.

It turns out that we always have extremely small error bars for the shear modulus and, typically, 10 times worse error bars for the compressional moduli. That has been carefully analyzed statistically on a lot of samples and based on how well we can determine chi squared for a Lee's squares fit on these elastic constants.

(Transparency)

This one is one of those things you are never supposed to show to anyone in a talk, but I cannot resist. This is six elastic constants and we fit 49 resonances to 0.16 percent. That is all you have to look at.

(Transparency)

How well do we really do? This is copper. These are literature averages of the elastic constants of single-crystal copper. These are our measurements. This is on a rectangular parallelepiped single crystal. This is a polycrystalline sample, rectangular and cylindrical. We did

the cylindrical problem. The cylindrical problem is nasty, because the mathematics does not work out so nicely.

Just look at the single-crystal numbers. This tells you globally how well people know elastic constants. I think the resonance measurements are the best.

MR. LAVRENTYEV: Why do you have six?

DR. MIGLIORI: Because this is wire-drawn. It has texturing and so it is hexagonal symmetry and it has five, and the last one is bulk modulus, and two of those are constrained.

(Transparency)

This is particularly interesting. This is tantalum. There was a big mystery with tantalum, because everybody measured elastic constants from room temperature down, and all of those guys who did not do that measured it from room temperature up. Right at room temperature there was a break in slope.

We made the measurements from about 100° K up to 350° K right through that region and, actually, there is a break in flow.

(Laughter)

In the process, we looked up all the best measurements of the elastic constants of tantalum and here are those guys, and here we are. You can see some people were not doing too well. But you can just look at the numbers for a second and see how well elastic constants are measured. What is very interesting about this one is that we did not bother to x-ray-orient the single-crystal tantalum crystallographic axes with the cubic axis, the axis of the sample, just for fun, because we wanted to see if we could find the microscopic orientation of the crystal lattice relative to the rectangular parallelepiped.

We did. We not only fit the elastic constants to what we believed to be better than all the other literature numbers, we found the orientation of the crystallographic axes in this tantalum single crystal. We checked it later with x-rays and they were within about one degree. So you can determine almost anything you would ever want to know.

(Transparency)

The way you determine it is computationally intensive. You need to know the resonances of a rectangular parallelepiped. The reason it works is there is lots of information here. These are some of the modes, low-lying modes, a very rectangular parallelepiped. There is not a plane wave among them. We cheer for that, because there is lots of information here. They are sufficiently complex to provide all of the elastic moduli and the mechanics problem on a computer is completely understood.

(Transparency)

This is the obligatory Rayleigh quote for any ultrasound talk. Trying to compute these resonances is an old problem and a difficult one. Rayleigh worked on this on his honeymoon. I have two quotes here and I will read them to you.

One from Rayleigh says that "In the case of a short rod and of a particle situated near the cylindrical boundary, this lateral motion would be comparable in magnitude with the longitudinal motion, and could not be overlooked without risk of considerable error."

What that tells you is that in short fat bars the mode spacing is not harmonic any more, it is not like the long rod, and it is not trivially found.

Then he tried to do a rectangular plate with free edges and basically failed with non-digital integrated circuit-type helpers around. This problem remained a fairly decent one for quite a while until Orson Anderson, who is really the father of this whole technique, and his student, Harold Demarest, and a few other guys figured out this wonderfully tricky Lagrangian minimization technique for solving the problem.

If you have a geometry that is easily described in a computer, then you can use some mathematical trickery to get computations that are as good as the measurement. That is unusual in resonance measurements. Usually you can check the codes by making measurements, not the other way around.

Finally, recently, meaning in the last 30 years, if you have a Pentium, or something like that, and you try using finite element codes (if you do not know what those are, I am not going to tell you), you wind up with a computational problem that goes like the cube of the accuracy with which you want to get the answer.

However, if you are smart, you can get this down to the square, which means that we can do computations on a P.C. in minutes. The heart of this whole business is in the computations. It is impossible to really go over this in any detail, so I am just going to outline it briefly. The trick here is not to do what the engineers do when they compute the modes of a bridge. They are really happy with three-digit precision and we can make a six-digit measurement.

(Transparency)

What, basically, Orson Anderson's postdoc figured out was that if you wrote down the Lagrangian for the solid

-- remember, it is based on elastic constants and densities and things like that, and we have a kinetic energy and a potential energy piece -- and you minimize it; that is, you allow the displacements of the oscillating object to vary arbitrarily, then --

(Transparency)

-- you wind up with a conversion to a surface integral, part of the piece comes out as a surface integral.

So you get two things, really. One of them is the wave equations, all of them, and the other is this very interesting little piece here, and this is what saves us. All we have to do is make this thing vanish and we have all the coefficients, which means, then, we know all the solutions and all the eigenvalues.

That is a very simple thing to do. I know I am just touching the surface of this, but there are a couple of interesting things.

(Transparency)

We have to make the surface traction, that is, any force that is on the surface be zero under those constraints. You can do it by making approximations to the displacements, u , and shoving them back into the integral until it works out okay.

It turns out you do not need to use an orthogonal set to describe those displacements, especially if you are using a computer. You can simply use x , y , and z to various powers to approximate the motions, pop it into the Lagrangian, you get this sort of relationship now. There are the eigenfrequencies we are after. We now have these matrices from kinetic and potential energy, and basically you are done.

You do that, you have matrices that you need to diagonalize, because there are lots of little pieces to the approximations to the motion and lots of pieces to the elastic constants. It turns out that to get six-digit accuracy, six-plus-digit accuracy for a frequency computation, you need to diagonalize roughly a 1000×1000 matrix.

However, if the system has symmetry of orthorhombic or better, the matrix block diagonalizes because of the separation of all those strains, and you need to do maybe only eight 100×100 matrices. Boom, 20 seconds on a Pentium, and you are done. Almost.

The trouble is that this is a prescription for calculating the resonant frequencies from the elastic constants. What we want are the elastic constants from the resonant frequencies. This is what we essentially spent two years studying in the computer: How to do this so that the codes would converge, would not get confused with different mode types, and also studying the error bars.

This has been written up and it is really better to read it than to listen to it, but it is a sneaky and difficult problem to satisfy yourself about, the systematic errors that creep in from the computation and how good your error bars really are. We have been very realistic about them.

We based them primarily on hundreds of pieces of data in which we believe that when we do a minimization, a fit to elastic constants, if the chi square changes by 2 percent, then we can compute an error bar for the elastic constants from that. That 2 percent seems to be about what we can do experimentally. The only reason we say that is because just time and again we can see if it is 2 percent off, we know what went wrong or we can find some other error, or whatever. It is a fuzzy error bar and those are probably the best kind.

(Transparency)

Using that code and that principle, it is important, because we have to use a complete set to describe the motion of the object, that the object shape be simple: rectangular, parallelepiped, sphere, cylinder, not too complicated. This is the potato. This is what we used to check the code on when we were running it on the Cray, before Pentium came out.

It is a sphere. It is really an ellipsoid with every semi-major axis different and cut in half. It turns out that object is no more difficult to compute than the resonances of a sphere. It can be machined to a part in 10^4 , so we found someone who had a numerically controlled machine designed to make bombs, who did not have much to do, fortunately.

(Laughter)

We had him make this object for us, with the promise that we would put his name on the paper, which we did. The code and the object agreed to the machining accuracy of the object. It is aluminum, it is 60-61, which has two peculiar properties. One is, the grain properties do not move around and another is that when you stress-relieve it, it never, ever textures itself, so it stays isotropic.

(Transparency)

This is the potato and we distorted a sphere in the ellipsoid direction, the potato direction, spheroid, sphere, egg, and so on. You can see the changes in frequencies. Of course, there are these degeneracy points. The sphere is a very degenerate object -- this mode and this mode are the same.

But when I make an ellipsoid, the frequencies shift and the degeneracies break. This is something we wrote up a while ago just to see what is going on.

(Transparency)

How do we make all these measurements? We want to measure the sample, not anything else, a standard principle of physics measurements. That means that we had to be tricky about the transducers. Diamond has the highest sound velocity of anything that you can hold in your hand -- not neutron stars and stuff like that.

In order for us people with fingers this big to make measurements on objects the size of the head of a pin, we really were not interested in handling and making transducer parts -- oh, the first rule of experimental physics? Make everything yourself. If you let somebody else do it for you, it will take six months and you will never get it done. Build it yourself.

We were going to build our own transducers. The idea was to use mostly diamond with a little bit of lithium niobate to act as the piezo element. That way we could build a large transducer with high resonant frequencies, we could use our codes to compute the response to the transducers. These transducers have a lowest mode of 4.3 MHz.

That means, by the way, that we could measure 100 resonances in a 2-mm sample of ordinary material, way more than we ever need to do any experiment.

(Transparency)

Here is a picture of the parts of the transducer, with a common object shown for scale.

(Laughter)

(Transparency)

Some of the transducers we epoxied together at low temperatures. We also developed techniques for putting together backloads. In this case, this is an alumina backload on a PZT-5A element (this was for some bigger work, but it is easier to see). We want this to act like one object. Epoxy does not do a great job of that.

So we developed techniques to silver-diffusion bond all the pieces together. This is a transducer that has no dissipative elements in it, no damping, no nothing. It is solid stuff, solid stuff, solid stuff, glued together with silver. The effect of that is twofold.

(Transparency)

One, epoxied-together transducers have resonances like this. All metal diffusion-bonded ones have resonances like that. Cool. The other thing is that the things that cause these resonances are goo, like epoxy. Epoxy is dissipative and the dissipation is a loss, just like a resistor. The loss makes noise, just like a resistor. The diffusion-bonded transducers are quieter. We can see that noise effect.

(Transparency)

We can also use some of those diffusion-bonded transducers at high temperatures. This is a resonant cell that goes from around room temperature up to 500° C. That is not in Orson Anderson's range, but it is so nice and convenient to use this thing and we do not have any buffer rods that we use it all the time. We can see lots of good physics with it.

We also have to prepare the samples really well. The geometry errors in the sample are the errors in our measurements.

(Transparency)

Once again, the crescent wrench strikes. There are lots of guys out there who want to sell you all these fancy pieces of hardware for aligning your crystalline and x-ray machine and grinding it and stuff. None of them are within an order of magnitude of the precision we needed to do our samples. We need at least a part in a thousand parallelism and perpendicularity for the samples.

Well, how do you do that? Do not buy anything. Go get a little glass plate, machine some shims of steel here --see them? -- trap the sample in a little pattern like that, and grind it. Flip the sample over, squish it this way, and you have now got this face perpendicular to that. Grind it, again. For 20 bucks we can do better than a part in 10^3 parallelism and perpendicularity in a little

while on the samples. We put a lot of time into these techniques but the answer is simple and does not cost anything.

(Transparency)

What can we see? Lots of stuff. Here is a Curie temperature, second-order phase transition in $\text{Nd}_2\text{Fe}_{14}\text{B}$ magnet. This is a single crystal. I am going to show you some stuff now and then we will get to physics, again.

(Transparency)

Curie temperature is the big magnetic effect in a magnetic material. It turns out that ultrasound is -- this is a very sensitive measurement technique. You can see very small changes in everything. In $\text{Nd}_2\text{Fe}_{14}\text{B}$, which is the magnet material, at room temperature the spins align along one of the crystallographic axes.

As you cool down, the atoms around shake less and less and leave more and more room in this cage for the magnetic moment to rotate. At temperatures around 140° K there is a spin reorientation and all of the magnetic moments rotate on to a body diagonal of the cell. It is a subtle physics effect, but it just jumps out at you when you make an ultrasound measurement. Magnetism and ultrasound are connected.

(Transparency)

We can do things at higher temperatures. Here is nickel. That is the Curie temperature for nickel with resonant ultrasound. Look at the numbers here. This is about a part-in-a-thousand change in the elastic constant. I am plotting frequency here, because it turns out that just as with the sphere, certain frequencies depend only on the shear modulus and other frequencies depend on a complex mush of c_{11} , c_{12} , c_{44} , but many of them are pure c_{44} .

When we get a frequency that is pure shear modulus, I do not have to do a computation. I need the computation to tell me which frequency it is, but once I know, I can plot shear modulus versus temperature by simply plotting the frequency versus temperature.

DR. GARRETT: When you are doing these temperature measurements, do you do what Dave and I do in the simpler case, do you choose a single resonance, phase lock on that resonance, and then track it with changing temperature, or do you do RUS all the time?

DR. MIGLIORI: I do RUS on all at every temperature, mostly.

DR. GARRETT: You do not phase lock on a particular resonance of interest?

DR. MIGLIORI: No.

DR. GARRETT: That is efficient for doing temperature sweeps?

DR. MIGLIORI: Yes, because I fit every Lorentzian resonance to get the center frequency and the Q at every temperature, so if I am trying to get elastic constants, I might measure 30 resonances at every temperature, fit every one of them with a Lorentzian, find the center frequency

and the Q, plot the Q for ultrasonic attenuation, use the frequencies to go into the code to get elastic constants, determine which resonances are which, and then I can track some of them later.

I might need only one, but I always sweep through them.

DR. GARRETT: The data you showed here with nickel was complete RUS on a complete sample, not on a single phase-locked mode?

DR. MIGLIORI: Right.

(Transparency)

We can see other second-order phase transitions. This is a discontinuity in the slope of the shear modulus at the superconducting transition in ytterbium nickel boron carbide. Again, check out this vertical scale, a part in 6000 here.

(Transparency)

We can help all the materials people. This is a valence transition. The valence of ytterbium in this metal changes at 40° K. Something else happened at 60° K. It turns out that something else was a mistake in preparation; it was a second phase creeping in, which later the people who made these materials got rid of, once they saw these ultrasound data.

(Transparency)

We can also see antiferromagnetism. Again, look at the vertical axis. This is the antiferromagnetic transition in a heavy fermion compound, uranium copper 5. Boom.

(Transparency)

But we can analyze some of these data really well. FeSi is a small-gap semiconductor. The origins of its semiconductor nature are associated with the classic physics problem, the avoided crossing. I have an energy band for one electron that looks like that, and another that looks like this.

Where they get near each other -- here is the Brillouin zone center, and this is $K = \pi/A$ out here -- I get an avoided crossing. This produces a little distortion here and here downwards, so I get a very tiny gap with almost perfectly symmetric properties above and below the gap, and huge density of electronic states near the gap edge. This is a theory.

Zack Fisk and I wanted to know if that made some sense and we made some single crystals of this sample and shipped them off to Stanford, where they did photoemission studies. Photoemission studies take electrons from a particle accelerator, our favorite physics tool, and they throw them

-- they must cost a dollar per particle --

(Laughter)

-- they throw them at these things and they look at the energy of the electrons coming off, they look at the energy shift.

The electrons come off the bands. If there are lots of bands a little bit off the Fermi energy, you see a big peak in the photoemission right off the Fermi energy. They did not see anything.

(Transparency)

So we said, gee, let us go measure the sound velocity, which we did. In fact, we measured everything and here is the elastic constants versus temperature. Remember, the elastic constants are related to the free energy. If we have those bands, I can write the electronic free energy very simply. I just sum over the energy of all of the electrons minus the chemical potential over kT and, boom, I have got it.

I kind of want to know the electronic free energy, so we used a system that does not show that narrow gap and measured the thermal expansion.

(Transparency)

So we took cobalt silicide and measured the thermal expansion and got its free energy, subtracted that from the model and --

(Transparency)

-- we were able to show that very large peaks in the electronic density of states right off the Fermi energy or the chemical potential produced the variation in the sound velocity that we measured in FeSi.

(Transparency)

Not only that, we could look at the ultrasonic attenuation of FeSi by looking at the width of the resonances and we find that the peak in the ultrasonic attenuation occurs at the activation energy between those two bands, sliced, diced, chopped, and peeled.

The poor photoemission guys went back and said, "Oh, we'll look at this more carefully," and then they found the density of the peaks and the density of states and everybody published in *PRL* and we were all happy.

(Transparency)

The point is that we took a classical simple physics measurement, utterly trivial, analyzed it carefully, took superb data with very high accuracy, generated a model, fit it, and out come details in the electronic structure of these metals. The reasons are that we understood this problem really well and we had very high accuracy and precision. Those will beat fancy techniques that are fuzzy and not very well understood many times.

(Transparency)

Here is another example: La_2CuO_4 . This is a high-temperature superconductor. Remember, we can measure all the elastic constants at once, because we are looking at a whole bunch of resonances. This one actually turns out to be tetragonal in its high-temperature phase, so it has six elastic constants. Two of those are shear moduli associated with edges. One of them crashes at 223°K . The other shear modulus does not do anything.

(Transparency)

We can do this same measurement and, in fact, we measured all the elastic constants of this material on higher temperature samples and we get the same dependence for the elastic modulus that changes, plus we measured the temperature dependences of all six elastic moduli.

This glitch here is an avoided crossing problem, where the edges of this sample are only about a part per thousand parallel, so they mixed some of the modes and there is an avoided crossing in this region that produced this glitch. There is an example of a really subtle systematic error that can make you crazy until you understand it.

(Transparency)

We can model this system. What is happening here is that La_2CuO_4 -- this is a little phase-transition stuff now

-- has a bunch of these oxygen octahedra. When things are hot, they are vibrating so much that they stand straight up. As things cool off, what happens is this octahedron, which has a square base, the base scissors like this and one axis gets too long and the other gets too short and they all tilt.

So they tilt like this, all the way down, and it becomes, instead of tetragonal with a square base, orthorhombic, with the axes shifted 45 degrees.

FACULTY: I think I understand the phase transition, but I am not sure I understand what you meant a while ago about a second-order phase transition.

DR. MIGLIORI: Nothing happens. What don't you understand about second-order phase transition?

FACULTY: I am not sure what it is.

DR. MIGLIORI: A second-order phase transition will be explained to you in about two minutes.

FACULTY: I know you paid me to ask the question, Al, but I was not sure when I was supposed to.

(Laughter)

DR. MIGLIORI: What is happening here is, as you cool down there is more room for everybody to rattle and they start to tilt to lower the free energy. At the temperature at which they begin to tilt, nothing happens, except, elastically, things become soft. It is now very easy to tilt them if they did not tilt, but they do not actually tilt, so nothing really happens.

We can define an order parameter for this system: the octahedron tilt. What does that mean? We are talking about a second-order phase transition.

(Transparency)

The way it works is this. At high temperatures I have a system in which the octahedra want to stand upright. At low temperatures I have a system in which they are slowly falling over,

because there is too much room in this lattice. The second-order phase transition is one in which that order parameter varies continuously as I cool; there is no sudden jump. Therefore, there is no latent heat associated with it, but there are discontinuities in the specific heat.

The free energy has a form like so. Now, let us look at this carefully. There are just some numbers here, α , β , c_{44} , and γ (c_{44} , of course, is going to be a shear modulus). There is a number, $T-T_c$, in front of the order parameter, Q , which is the angle at which these octahedra tilt.

If I plot this function, and T is greater than T_c and α , β , and γ are positive, I get this curve here. If $T-T_c$ becomes negative, I get this curve here. So if the order parameter is the minimum in the free energy, it stays at zero until all of a sudden T goes below T_c , and then it starts to move out from zero to this minimum here.

That is a kind of generic, what is called Landau-Ginzberg, description of a second-order phase transition.

What I can do with this, then, is to predict stuff. For example, if I have got two octahedra tilting like this, what they are going to do is take a square and turn it into a diamond. That looks just like shearing it. So I know that I can connect one of the shear moduli to that tilt, and the reason I know it is I think about it. There is no little lab view for second-order phase transitions that tells you what the order parameter is going to be. You have got to think about it.

I know what it is going to be, and so I put it into a little free energy equation. This stuff here is what produces this quartic that has two minima here and one minima there.

This is the ordinary elastic energy of the solid. If I put in a shear strain, e_{44} , and I have a shear modulus, c_{44} , the elastic energy is $\frac{1}{2} k_x$. Then I couple the strain to the order parameter. In this case I am going to couple it linearly -- here is strain, here is order parameter. If the octahedra tilt and it goes from a square to a diamond, then there must be some coupling between shear and the octahedra tilt, meaning if I shear it, I am going to probably force them to tilt.

Throw it in and see what happens. What I am going to do with this now is I am going to compute the elastic response under the constraint that the free energy is a minimum. This is the whole way of using elastic constants to get information about the phase transition.

(Transparency)

The shear modulus is the strain susceptibility of the free energy. $D^2 F / D \text{ strain}^2$ under some constraint if the strain is zero. I am not going to go too carefully into that, because that will take us some bit far afield. What I do is I take the derivative of stuff and set it equal to zero. This is the thing I use to find the order parameter. That is the derivative of the free energy set equal to zero with respect to the order parameter.

I take the derivative with respect to strain, set it equal to zero, and I compute the derivative of the order parameter with respect to strain under these constraints. I can then go ahead and take the

second derivative -- you can work this out later, it is kind of fun -- to get the elastic constants, put in all the constraints, and I get some very interesting things that happen.

If T is greater than T_c , the solutions to the constraint and these problems tell me that c_{44} goes like the undisturbed elastic constant minus something over $T - T_c$, so it is going like this. Below T_c I find that it also goes like this, but with a little factor of 2 in here. What I expect, if I have that linear coupling between strain and order parameters, is one of the elastic constants to go like that.

I could change the coupling somewhat differently. Look, it is really hard to believe that you have linear coupling, however. I am setting you up for this. Why should it matter whether I shear it this way or this way which way the octahedra tilt? Remember, they are forming a diamond. I am going from a square to a diamond shape by tilting things at the corners.

It cannot possibly matter which way they tilt, because if this one tilts to the left and this to the right, I get this diamond, but if this tilted to the right and this to the left, I am still going to get that diamond, because this is really a crystal lattice. So it makes absolutely no sense to have linearly coupled the strain to the order parameter. That is a key. I just use a symmetry of the crystal and a symmetry of the physical effect to tell me something about how the order parameter had to couple to the strain.

This you can proceed with *ad infinitum*, or *ad nauseam*, piece by piece. But if you have a model for the coupling, you can tell what all the elastic constants are going to do, and if you have resonant ultrasound you can measure them and check the model. And that is what we do.

It did not make any sense to do that linear coupling, but it did fit the data, as it turns out.

(Transparency)

We could do quadratic coupling. That is where it does matter which way they tilt. So now I have a Q_- over here instead of a Q . I go through the same analysis, dF/dQ , set it equal to zero to get the parameter. I note that Q is zero above T_c , just the way we expected. I can now compute what Q_- is below T_c , put it back into everything, and now I find the following.

If T is above T_c , c_{44} is c_{44} . If T is below T_c , c_{44} is c_{44} minus something. There is step discontinuity in the shear modulus. That is the physically reasonable thing to find, but it does not fit the data. This is where the puzzle came in.

We then went and measured all of the elastic constants and we were able to show that no simple analysis like this whatsoever will produce the measured elastic response of La_2CuO_4 --

(Transparency)

-- unlike that paper I showed you at the very beginning, where this sort of analysis was done because they had data on only one mode, it looked linearly coupled, they threw it in, did a Ginzberg-Landau, and published it. That defined what was going on.

That is not what happens. In fact, you have to use a model in which the octahedra do not tilt, they buzz out into a quadripolar "buzzy" shape to get the correct elastic response. It is wonderful.

Again, we got microscopic details of the motion of these structures inside a high-temperature superconductor, one of the very interesting problems of the last few years. We got at it with a simple, well-analyzed resonance measurement -- a lesson to me, for sure.

(Transparency)

We have got at other physics as well. This is a measurement of deuterium and yttrium. What we have measured here is $1/Q$. This is a metal that absorb hydrogen like crazy. The hydrogen goes in interstitially. It can move around, it has some mobility. That mobility, when you compress the material, you change where the hydrogen is, it generates entropy, because the hydrogen pops one place and then another, and every time it moves you lose all the energy of the configuration.

You can measure the hydrogen motion and some details of it and model it according to a Debye relaxation model, which Dr. Bass talked about a little bit. We can measure $1/Q$, find the peaks in the ultrasonic attenuation as a function of temperature, which tell you the temperature that matches up best with the energies of motion of these deuteriums.

Then you can do it for hydrogen-deuterium and you come up finally with a model that tells you in pretty great detail how the hydrogen is stuck in there and what it is doing.

(Transparency)

The last bit of stuff. This is a Trident II missile.

(Laughter)

This is my favorite picture ever taken. This is like \$100 million worth of Dr. Hargrove's hardware made by Dr. Hargrove's buddies. This is the first launch of this thing from a submarine. So it got up -- swish -- went out of control and they blew it up about right there. The reason is that there was a quality control problem.

(Transparency)

Resonant ultrasound can be used to help with quality control problems. The way it works is that it is very sensitive to the elastic constants but, also, the geometry of the object. This is a tiny single crystal about the size of the head of a pin of Europium hexaboride, which is a material that has -- it is used for filaments in electron microscopes because its work function is lower than anything, which means that you can heat the filament less, which means the electrons come off colder, which means that the noise they produce in the images is lower.

This is a little sample and we were measuring resonances of it that connect to its free energy. Something happened to it and the resonances went like this. We later looked at it very carefully and we realized a 100-Å crack had developed in it and the crack was buzzing. It killed the Q and a lot of the resonances and reduced them, so we could clearly see a crack.

DR. GARRETT: That spectrum, then, would be amplitude-dependent?

DR. MIGLIORI: Yes, and of course, all the modes are amplitude-dependent the way we do them, because we have no control over exactly how we excite or where, so amplitude information is really very bad to use, unless you are desperate or you are doing nondestructive testing and you are an engineer and you do not care.

(Transparency)

We can see cracks because it changes the Q and it shifts the frequencies. We can also see changes in symmetry. This is a cylindrical rolling element in a cylindrical roller bearing. There is a burn mark on this one. It measures perfectly cylindrical, but because they burned it in grinding, there was a soft spot that shifted the elastic moduli very slightly in the steel in a noncylindrically symmetric way, broke the degeneracy of this resonance into two peaks, like this.

So you can see something that you could not measure, but you can see a heat-treating error.

(Transparency)

You can also clearly see errors in the sphericity of ball bearings. Here is a bad silicon nitride ball bearing. This is a split mode. We know exactly the degeneracy of every mode. We can compute how many there should be when they split. From the splitting and frequency we can compute how out of round it is, if it is not textured; that is, if the material it is made of is elastically isotropic.

Unfortunately, in silicon nitride it is a hexagonal basic material and occasionally manufacturers accidentally texture it, so that this splitting, in fact, is from a perfectly round ball that is elastically anisotropic. It was ground round (that is not food) at room temperature and it was very round, but not elastically.

Well, not only is it not elastically round, its thermal expansion coefficient is not isotropic. So when this ball is taken up to 400° C, where it is supposed to be operating, it goes thunk, thunk, thunk, because it is not round any more and we can make an elastic measurement at room temperature that tells you the out of roundness at 400° C.

(Transparency)

All of this makes people who want to make money making instruments interested in us, unfortunately. After huge amounts of time explaining how resonant ultrasound worked to a bunch of people in Albuquerque, they built this machine here, which is the big brother of the one on the back bench. You guys can put seashells and beer cans in that later and play with it -- I am going to leave it there all day.

General Motors makes 45 million zirconia, which is zirconium oxide, oxygen sensors for their cars every year. They cost them about nine cents apiece, the raw ceramic part. However, when one of these things fails in your Chevy, the little light comes on, on the dash, and it says, "Go buy a Japanese car."

(Laughter)

So they do not want these things to fail. It costs them 150 bucks to replace them at that point.

The way they checked for flaws, they took these things and they pressurized them until they broke or not. They took them up to some pressure, if they broke they did not put them in the cars. However, they damaged a lot of them in the process. The success of that particular test was about 70 percent, meaning they still let through about 30 percent flawed parts.

If you look at the resonances, however, you can find cracks.

(Transparency)

This is the separation between two nominally degenerate modes. This is one flaw they were looking at. This is another flaw they were looking at. You can just barely do it. It turns out that "just barely" is roughly the following.

(Transparency)

This is a rundown of a statistical test on 56,000 pieces, of which known good and bad ones were put into the system and run through. You went from 10 to 40 percent flawed parts to 0.3 to 1 percent flawed parts making it through, and your reject rate on good parts was not bad. That is not too bad. There was not any easy test, and you can do the test in one second.

(Transparency)

Then, of course, we did do the highway bridge. This is the dumbest -- no, the second dumbest -- thing I ever did.

(Laughter)

Check this scale. You could stand on this bridge and feel it going at 2 Hz. Before we cut the I beam we measured these modes by putting -- we took a truck and we had this hydraulic cylinder about this big around that shook up and down with a hydraulic motor and produced 2000 pounds of force on the bridge at 2 Hz. That was our drive.

We could measure the resonances, and then we cut the I beam, and they shifted. You can extrapolate back, but I do not think that technique will go anywhere, ever.

(Transparency)

I will quit with the quote I keep muttering to myself all the time about physics, and the little lesson in nature here. Do you all see the insect in that picture below there?

See, you look at this nice picture and there is the bug: It is a walking stick hidden in there.

There are lots of these things hidden in physics, I hope. Thank you.

MS. LU: Are you able to determine residual stresses by looking at frequencies?

DR. MIGLIORI: The answer is yes, under very favorable circumstances, meaning extremely good geometry parts, where you already know the variation of elastic moduli with stress. For example, in steel it is quite low. You can put a lot of stress on a good steel before you change the

moduli very much, and it all depends on details of, basically, the pressure dependence of the elastic constants.

MR. LAVRENTYEV: If you have an isotropic cube, do you have to know the orientation of the axes?

DR. MIGLIORI: If we know the orientation of the axes, the error bars go down and the computational problem is much easier to solve. If we do not know the orientation of the axes, we can compute them and fit them. But we need more data and they have to be better.

MR. LAVRENTYEV: If you have high superconducting materials, how do you machine them?

DR. MIGLIORI: Very carefully. The technique I showed you, where we have these metal shims that completely surround and support the sample prevent the edges from crumbling as well when you grind it. There are all sorts of benefits to that, so we can machine them. We have the only full set of elastic constant data, I believe, on 1-2-3 as well as 2-1-4 samples; all the moduli are accurately done.

Is that right, Moises?

DR. LEVY: That is right.

STUDENT: In the magnet I found the frequency about 100 -- the one slide you showed.

DR. MIGLIORI: No, I do not think I showed any data today on our measurements that were above about 2 or 3 MHz. If it is, it is a typo.

STUDENT: How....the pores in the surface?

DR. MIGLIORI: You do not need an optical polish and if you remember that we are approximating the motions with roughly 10th order polynomials to get roughly six-digit accuracy in the modes, you can see that a surface that looks matte to your eye, not perfectly polished, will work.

DR. MARSTON: You had a picture here about rocking the sample to detect --

DR. MIGLIORI: Oh, yes.

DR. MARSTON: Can you say something about that?

DR. MIGLIORI: One of the puzzles was how do we identify the modes. We thought that we needed to identify the modes before we could do a computation. Every mode has a different motion and, in fact, because we could do the computational problem, we knew how the corners were supposed to move.

Many of the corners move almost perpendicular to the body diagonal, within about 10 degrees of it. So here is the body diagonal of a cube sitting up in our apparatus and that corner is moving like this. If I really think I know what I am doing, I can tilt that sample until the corner motion is parallel to the surface of the transducer and then I will get no signal.

What we did was to build an apparatus where we could run a stepping motor and tilt the sample all over the place and look for the nulls in the signals at each mode frequency. We then could plot the directions in real space of those nulls, and then compute them with the code and we got perfectly good agreement.

At that point we said the code is working, we understand the modes, we do not need to know what they are to get the computations to work, the problem is done. That was really the last thing that convinced us that we really understood all the mechanics that were going on.

DR. KEOLIAN: I did not understand how the trick with the force steel blunts grinding the sample on the glass plate gave you parallels.

DR. MIGLIORI: First, I take a sample that looks like a piece of rock I got out of the garden, plop it down on a glass plate, grind the surface, and now I have got a flat. Flip it upside down, put it on the glass plate, grind it, again. I now have a flat that is very parallel to the first one, because those shims are all of equal thickness.

I just ground a flat. Now I flip it over, put that flat straight down. It is touching the glass. The shims hold the sandpaper parallel. I cut another face. Now I have two parallel faces and a lot of garbage surrounding the rock.

DR. KEOLIAN: So you are grinding on both sides?

DR. MIGLIORI: Yes. Now I take the sample out --

DR. ATCHLEY: What kind of sandpaper do you use?

DR. MIGLIORI: It is the paper they use to polish fiberoptic cable ends, so it is 0.3 μ grit.

DR. SWIFT: And it does not take the steel down, it just takes the sample down?

DR. MIGLIORI: The steel area is 10^4 times the sample area, so it does not take very long and every once in a while you re-machine the shims. The measured error bars are less than a part in a thousand routinely.

DR. HARGROVE: But the results are only as good as the shims?

DR. MIGLIORI: Yes. Okay, so now I have got them like this, right, two parallel faces, and

--

DR. LEVY: How do you attach the shims to the sample?

DR. MIGLIORI: I use a molten wax to hold them to the glass plate so I can manipulate them and hold everything down with a magnet underneath while the wax is cooling. I sandpaper it, melt it, flip it over, sandpaper it, and then I take the two parallel faces and I push the shims sideways together to grab the parallel faces instead of the bottom one. Now I have got this one perpendicular to these two. Flip it over, perpendicular, turn it, flip it, boom, you are done.

DR. KEOLIAN: You are just taking the parallelism that you have somehow in these shims and bringing it over --

DR. MIGLIORI: Yes, the Bridgeport-Allen machine can do a part in 10^5 errors in angles in machining the shims without the machinist really looking at the tool very well.

The real bottom line is we tried it, checked it with very accurate metrology at Los Alamos and it worked, and then we quit.

DR. GARRETT: Could you say a few things about your preamp design for those of us -- Dr. Hofler seconds that request.

DR. MIGLIORI: My time is up. I will be glad to do that, say more than a few words about it and discuss it. Would you like to do that in the evening one of these times? I will show you the detection scheme and the rest of it, too.

DR. BASS: We do have time for maybe one or two more questions.

DR. HOFER: Did I miss how you measured the temperature of these samples? Obviously, you do not want to stick something onto the sample.

DR. MIGLIORI: Right, and I have very poor thermal contact through the edges, so I have to surround it with helium gas. That is easy, because mostly I am in a low-temperature environment. I pump the gas out to where the mean free path is plenty long, so I might use 1/10 mm of gas pressure. I have extremely good thermal contact.

We can check it on many of the samples that have sharp phase transitions by going up and down and looking for hysteresis. It is not as good as the world's best cryostats, but we will typically have errors of about 10 mK between thermometer and sample and temperature flux rates. I do not care at that level; it is not important.

MR. WANG: Do you have some kind of criteria by which to judge if your resonance peak is good by looking at only the shape of the peak or do you say based on the Q of the peak to say this resonant peak is good?

DR. MIGLIORI: I know what kind of Q's I need to get the precision I need in advance for whatever the particular experiment is. I can fit the peaks to determine if they are Lorentzian or not. I know how to add in coherent background for minor errors in R.F. shielding. I understand the problem. If they look weird, then we do not use the data.

The definition of weird depends on the experiment and what we are trying to find out. So yes, you have to be careful. You can certainly publish garbage if you are careless.

DR. MARSTON: You mentioned early on that -- of course, phase is a very good way to infer a mass-and-spring frequency to remove dissipative information, yet all of the fits you talked about are fitting lines, so do you ever use phase in any of your data?

DR. MIGLIORI: Never. The reason is I am after every last bit of signal to noise I can get. The only way to get all the signal to noise is to use all the data. The only way to use all the data is to use every single point in those resonances to fit the resonant frequency.

If I try to find out where the phase crosses zero, I am using only a few data points. The effect is to throw away signal and increase the noise. If you think about it, sit down, you will see that that is what happens.

DR. MARSTON: But presumably there are going to be physics situations where you want to get rid of that dissipation and then you will have to use phase information.

DR. MIGLIORI: Right, but I am going to fit it, then. I am going to use all of it through the whole sweep.

DR. SWIFT: So your receive electronics could be looking at the real and imaginary parts of the signal coming in rather than just the magnitude. You have two functions to fit --

DR. MIGLIORI: Oh, that is pretty straightforward. If you take a Heterodyne detector, where you are trying to get rid of $1/F$ noise by using an I.F. frequency that is not zero, then you have to keep track of the phase of the I.F. frequency to get the phase information from the signal. That means another mixer, a second mixer, where you mix the drive with the I.F. to get the phase of the 1 kHz out. It is perfectly easy to do.

Or you go to DC Homodyne detect by using an I.F. frequency of zero. That picks up all the $1/F$ noise in the electronics and your noise is degraded. We made a decision quite a while ago not to use the phase information, which is not very important if the Q 's are as high as the Q 's that we show you in determining anything.

The factor of square to 2 that we throw away in signal to noise by not collecting real and imaginary parts is overwhelmed by the reduction in $1/F$ noise and other problems in doing DC. It is just a choice.

PERIODIC, RANDOM, AND QUASIPERIODIC MEDIA

JAY MAYNARD

PENNSYLVANIA STATE UNIVERSITY

DR. MAYNARD: I want to encourage students to ask questions during my talk, because if you do that, I am likely to get confused, accidentally skip whole parts of the talk, and we can play more volleyball.

I would also like to make a comment about this paper that was sent to the students ahead of time. If you look at this paper, it has got these horrible equations in it, lots of details, and you are thinking, "This is going to be a dreadful talk." Well, that may be right, but it will not be because the talk follows this paper.

This paper serves two purposes. One is that in the back it has the references to papers that do follow the talk. The other thing is that the talk is about waves in periodic, random, and quasi-crystalline media. Certainly, for periodic and random arrays of scatterers, you might run into that in your future careers. If you do, you may have to do some serious calculations for wave propagation in periodic or disordered arrays of scatterers. In that case, you are going to want this paper. It tells you all the details about how to do the calculations.

The reason I wrote this is that I found that those calculations are actually not in the literature, which is kind of surprising. They are straightforward, but they do get pretty messy and maybe that is why they are not published. But it is nice to see the details at some point. But right now it is not published, so this is not available in any store, but I am going to be submitting it soon -- I finally got it put together.

DR. GARRETT: Where will it be published?

DR. MAYNARD: In *JASA*. I am going to try to make it a tutorial article.

Dr. Garrett said that this is going to be the day of solid state physics. Actually, it turns out I am going to start with solid state, but I will get to acoustics soon enough.

(Transparency)

Here is an integrated circuit. In fact, it is a CPU. Actually, if you do computer simulations of acoustics, maybe you would think this is acoustics and not solid state. Anyhow, it is an integrated circuit, and of course you want to make it go faster and faster.

One way to do that is to make it smaller. If you make it smaller, you have to ask if the physics for the larger transistors and larger wires still holds for the smaller scale. It turns out that if you take the connecting wires in an integrated circuit, they may be only a micron or 100 Å wide.

You find that under certain conditions there is a dramatic deviation in Ohm's law for a wire that is this size. That is what we want to understand in solid state physics.

(Transparency)

That goes back to a very fundamental problem in solid state physics, and that is to solve the Schrödinger wave equation for an electron moving in the potential field of some 10^{23} ions or scattering sites. Now, 10^{23} is a big number. In fact, this number here -- at least in California -- is known as the "avocado" number.

(Laughter)

If you tried to use Mathematica to solve this wave equation with that many boundary conditions in it, it would take quite a while. But fortunately some solids are crystalline, so that the ions are arranged periodically and in that case you only have to solve the Schrödinger wave equation for one unit cell that may contain only a few atoms, and that is something you can do. I do not know about Mathematica, but electrical engineers have computer programs that they use for calculating the band structure of semiconductors and so on. So you can do that.

But unfortunately, in the real world a lot of solids are not perfect crystals. They can have disorder. In fact, there are some important material, like amorphous materials or alloys, that are highly disordered. Now, if it is highly disordered, you are back to having to solve the Schrödinger wave equation for some 10^{23} nonperiodic scatterers -- a big problem.

Fortunately, in the real world measurements are done at finite temperatures. At finite temperatures in a solid those ions can be moving; in other words, at finite temperatures a solid has phonons running through it. An electron can inelastically scatter off of a phonon. This inelastic scattering actually destroys the wave nature of the electron.

It randomizes things, and because of this randomizing by the inelastic scattering you can use statistical physics and convert the wave equation with some 10^{23} parameters into a transport equation, like a Boltzmann equation or a diffusion equation, that may have only a few parameters, like an effective mass for the electron and a mean free path for that inelastic scattering.

These are equations with just a couple of parameters that you can solve with Mathematica. Because of this inelastic scattering and the use of statistical physics, people have been able to handle highly disordered materials in solid state physics.

Two things happened that changed all that. One thing that happened is that the experimentalists made their samples so small and so cold that the size of the sample was smaller than a mean free path for inelastic scattering. So without the inelastic scattering, things no longer average to zero. Statistics is no good. You are back to having to solve the wave equation for some 10^{23} nonperiodic scatterers. Big problem.

This is actually a very fundamental problem. It is just wave propagation in a disordered array of scatterers. That problem dates back to Lord Rayleigh. Lord Rayleigh worked on it and since Lord Rayleigh a lot of famous people have worked on it. Some progress was made, but when the

problem came up with these small cold samples, a lot of people in condensed matter physics got to work on the problem and they actually made some interesting progress.

These contemporary theoreticians did not want to admit that they were working on a problem that Lord Rayleigh worked on, namely, solving the wave equation. They said they were working on systems that have long-range phase coherence. This has been a big frontier area in condensed matter physics, with a lot of buzz words like "Anderson localization" (which I will be talking about today); "universal conductance fluctuations"; "Aronov-Bohm effect"; "normal electron persistent currents"; lots of *Physics Review Letters* generated out of this.

There is also something called superdiffusion and that is something that I will be trying to talk about at the very end of this talk.

These theoreticians, the way they handle the problem was, in fact, not actually trying to solve the wave equation, but they go back to the equation that they can solve and then do perturbation theory back toward the original wave equation. They have these quasi-particles called diffusons and they try to include as much phase coherence as they can. Without the inelastic scattering, the electron does now maintain phase coherence when it propagates through the sample. So these small cold samples are one thing.

The other thing that happened that changed things is that quasi-crystals were discovered. These quasi-crystals are an aluminum alloy made by rapid quenching. Originally, the samples that you got were just very small and would be $100\ \mu$ crystallites -- actually, quasi-crystallites. I will be saying a lot more about quasi-crystals later in the talk.

(Transparency)

Let me go back to the experiment that people did with these small cold samples. They made one-dimensional wires that were about $100\ \text{\AA}$ across and $1\ \mu$ long and they measured electrical conductivity in these wires and found that it did not obey Ohm's law.

They gave a name to this kind of physics, this scale of physics, and it was called "mesoscopic" physics. It is looking at phase coherence on the scale of microns. For our research, we decided to coin a phrase, so we called it "megascopic" physics. That is physics on the scale of millions of microns.

But of course a million microns is a meter. So we like to claim that we measure phase coherence not in a $1\text{-}\mu\text{m}$ wire, but in a one-dimensional wire that is more than 10-m long. We study the density of states, etc. in a quasi-crystal. It is not $100\ \mu$ in size, but it is more than a meter in size -- easily the world's largest quasi-crystal.

Both of these statements are worded so as to shock the people who actually work in these areas. Our one-dimensional wire is not one of these $100\ \text{\AA}$ wires that they use for measuring conductivity. Our wire is actually a piece of steel wire, music wire, that we bought at the music

store. Our quasi-crystal is made out of aluminum alloy, but it is the kind of aluminum alloy you find in a machine shop.

Our systems are actually classical analog systems and, of course, since I am here, they are acoustic analog systems. One way you can think of these analog systems is as if they are analog computers.

These analog systems have some very nice advantages over the original solid state or quantum mechanical systems. That is, first of all, they are very precise analogs of the quantum mechanical systems. Let me go into that in a little more detail.

(Transparency)

In quantum mechanics you have the Schrödinger equation. The quantum mechanical systems are special, because you have gotten rid of the inelastic scattering and, in that case, you are dealing with the time-independent Schrödinger equation. In that case, you can assume that the wave function goes $e^{-iEt/\hbar}$, energy over \hbar times t . You plug that in. You divide through by $-\hbar^2 \nabla^2 / 2m$ and you get this equation here.

So this is an eigenvalue problem. This q here you can think of as an eigenvalue parameter. In quantum mechanics it is just the square root of $2mE$ over \hbar squared, and the E here are the energy eigenvalues.

You solve this thing for eigenfunctions, which would be the electron wave functions, but you can get solutions for only certain values of the eigenvalues, so you get discrete energy levels and so on, and that is the basis of quantum mechanics. The point is that you can write the equation this way.

In acoustics, we have the wave equation. We can always Fourier analyze in time or, in other words, just assume a time dependence, $e^{-i\omega t}$. But then if you construct an acoustic system where the speed of sound varies with position, either by changing the mass per unit length or changing the mass or somehow changing the tension, you can get a speed of sound that depends on the position, plug that in, and you can convert that wave equation to an equation that is mathematically identical to the quantum mechanical equation.

In acoustics, the eigenfunctions are the normal modes of vibration and the q now would be ω / C_0 , where ω represents the eigenfrequencies, otherwise known as the natural frequencies, or the resonant frequencies if you drive the system.

These are eigenvalue problems. The whole idea is to find the eigenfunctions and the eigenvalues. They are mathematically identical. All of the physics in the system is contained in the potential field. The key part of the potential field that we are interested in is its symmetry, whether it is periodic, random, or quasi-crystalline.

We can rig up acoustic systems that tell us the same thing that we would learn in the quantum mechanical system.

(Transparency)

But there are other advantages. These are megascopic systems, a meter on a side. In such a large system, all the conditions and the parameters can be very precisely controlled, or at least measured. If you are trying to polish a 100 Å wire, that is quite difficult. It is much easier in one of these megascopic systems.

Also, in a measurement in our analog systems we can get very precise measurements of the eigenvalues, the eigenfrequencies, and the eigenfunctions, density of states, and so on, all the same sorts of things you can try to measure in solid state physics.

In our system we can measure the whole eigenfunction, amplitude and phase. That is something that is virtually impossible to do in quantum mechanics. In fact, in these eigenvalue systems, once you measure the eigenvalues and the eigenfunctions, you are finished. Everything else can be calculated from those.

These first three things you can also do with digital computer simulations, and people do that, but it turns out that with these analog systems you can do more. You can make this potential field not only depend on space and have some symmetry, but you can also make it time-dependent. Also, it may turn out that you have a system where the forces in the system depend on the actual amplitude of the wave function, so you can make the thing nonlinear. In our acoustic analog systems you literally just turn up the volume. So it is quite easy to measure time-dependent and nonlinear effects in these analog systems. Those are things that are extremely difficult to do in computer simulations, very difficult. So we can learn new things.

Before I tell you about the actual experiments that we did on this megascopic physics, I want to explain how we got into the business, because it involves some transition between solid state physics and acoustics. It is not as straightforward as the resonant ultrasound.

At Penn State there is a big acoustics program, where they study many different kinds of acoustics, not just physical acoustics.

(Transparency)

Another area of acoustics is noise control. Here is a typical problem in noise control. You have a plate, like the floorboard in the little aircraft that fly into Monterey -- or, in my case, that do not fly into Monterey. At one end of this plate you might have a source of vibration that puts transverse waves in this plate. Those transverse waves travel down the plate. As the plate vibrates, it pushes the air, radiates sound, and makes some annoyance.

Usually, for structural reasons, a plate will have a rib on it. That rib reflects the vibrations, so there is less vibration transmitted, so there is less noise at the other end of the plate.

Up north, in Seattle, they have the Boeing Aircraft Company. The engineers at Boeing said, "Well, if one rib reflects the vibrations, why not a whole bunch of ribs?"

(Transparency)

So either for ease of manufacture or aesthetic reasons, they put down a nice periodic array of identical ribs. They found that for some frequencies the first couple of ribs will reflect the vibrations and all the rest transmit without any further reduction.

This is something that, if you have had a solid state course, you learn about. It is the same situation for an electron in a metallic crystal. In a metallic crystal you have electrons and positive ions. An electron is very strongly scattered by a positive ion (in fact, that is the strongest force that there is).

Now, if the electron is strongly scattered by the positive ion and you try to calculate electrical conductivity in a metal, you would have to conclude that a metal should have very poor electrical conductivity. The electron, in trying to go from one end to the other, just keeps getting scattered out of its way by these positive ions.

But of course, metals have high electrical conductivity. There are two reasons. One is that the electron behaves like a wave (that is quantum mechanics, again). The other reason is because these positive ions in a crystal are arranged periodically.

The analogous thing here is that you have a periodic rib structure or a periodic array of ions or periodic scattering sites, then a wave can go through the system as though there were no scattering at all. That is why metal can have a high electrical conductivity.

(Transparency)

Both of these things are consequences of a theorem; in mathematics it is known as Floquet's theorem. Of course, in physics, we had to rename it, so we call it Bloch's theorem. What Bloch's theorem says is that if you have a system with a periodic potential or, in acoustics, a periodic impedance, then the eigenfunctions for that system are extended.

In particular, the eigenfunctions have to have this form. There is just a phase factor, e^{ikx} , that almost, unfortunately, looks like a plane wave but is a phase factor, and then it's times a function that is periodic with the period of the potential field.

If you take the modulus of this wave function, then the phase thing just goes away and all you are left with is this periodic part. Because it is periodic, it looks the same in every unit cell. If you take a look at some sort of nominal or average amplitude of this wave function, it looks the same in every unit cell. For an infinite crystal, you get the same nominal amplitude everywhere. That is an extended eigenfunction.

I have taken a textbook -- this is actually Tinkum's group theory book (my third most popular physics book). In this one book on group theory there is a lot of physics explained clearly.

Anyway, these are examples of extended Bloch eigenfunctions,. This is the periodic part here, but you throw in the phase and you get some change, like this. You will notice that if you

took the amplitude of this, you would still get the same nominal amplitude, and this continues on in an infinite system. Here, the same nominal amplitude.

There is some structure in there in each unit cell. For the infinite system it still has the same nominal amplitude.

Bloch's theorem and the existence of the extended eigenfunctions are pretty much taken for granted in solid state physics. But if you actually try to do calculations, you will find it is a lot more complicated than the textbooks would tend have you think. In fact, that is a lot of what this paper is doing, some of the really messy calculations to get the actual quantitative results (rather than just the qualitative things).

By the way, for an electron in a metal, if you have this extended eigenfunction, if the electron wave function has the same nominal amplitude everywhere, that means the electron has equal probability of being anywhere in the system or, in other words, it can travel freely through a periodic system.

(Transparency)

This viewgraph restates the situation. If you have a periodic potential field, then the eigenfunctions are extended. That is sort of it for wave propagation and periodic potential fields. Bloch's theorem and extended eigenfunctions are well known -- I would not say they were well understood, having done some of these calculations now -- but at least they are well known, and you can find how it works in most textbooks and you can reproduce my calculations or just look to see how I did my calculations.

But now suppose you take a periodic potential field and you add a part to it that is random. What you might think would happen is if you add a little random part here, then this extended eigenfunction would just get little random variations in the amplitude of the eigenfunction. You might think that we go from periodic to a random part and we might just get random variations in the amplitude.

It turns out that that is not what happens. What happens is the eigenfunction, when you put in this disorder, becomes exponentially localized; that is, somewhere in the system there will be some point where the eigenfunction has a maximum amplitude. Then, as you go away from that site, the amplitude of the eigenfunction decays exponentially, and there is some characteristic length to this exponential decay. That is called the localization length.

This is not an obvious thing. I would have bet money that putting in random disorder would have put random fluctuations into an otherwise extended eigenfunction. So why do you get this exponential decay?

It is not a trivial matter. Phil Anderson and Sir Nevill Mott used this effect to explain the metal-to-insulator transition in disordered solids. For this, they won the Nobel Prize, so this is a pretty significant thing we are dealing with here. In fact, it is called "Anderson localization."

I have said that wave propagation in a disordered potential field is a very fundamental problem -- it dates back to Lord Rayleigh. I have been told that Lord Rayleigh almost got Anderson localization, but I guess he did not have Mathematica, so he did not quite get it, so he had to win his Nobel Prize for a different reason.

Anyhow, it is a quite difficult problem. Since Lord Rayleigh, a lot of famous people worked on the problem: von Neumann, Wigner, Freeman Dyson. They actually made a lot of progress.

The first rigorous theorem for waves in disordered potential fields comes from Furstenberg's theorem. Furstenberg was a mathematician -- among mathematicians this is a famous theorem (like Floquet's theorem). Furstenberg's theorem was in 1963, which is quite recent on the time scale of Lord Rayleigh. It is a fundamental problem, it has been around for a long time, and it has been solved only quite recently. And it has been solved in only one dimension.

In two and three dimensions, for disordered scattering fields in two and three dimensions there are no rigorous theorems. There are arguments, so-called scaling arguments, but there are no rigorous theorems. It is a very difficult thing.

When I first heard about Anderson localization, I was really puzzled, because, as I said, I would have bet money that you just get random variations in the amplitude. How does this disorder give you this exponential localization of the wave function?

I wanted to understand how that happened, so I went to the literature. In the literature there are thousands of papers published on Anderson localization, but all of those papers use Anderson localization as a way of explaining some observation in solid state physics. None of the papers gave a balls-and-springs explanation of why disorder gives you exponential localization.

I tried to read Furstenberg's paper. Furstenberg is a mathematician. His paper is 25 pages long and it is filled with theorems and lemmas. And it references other papers that are filled with theorems and lemmas. I am an experimentalist and I have a lot of trouble reading papers that start off with "Let M be a manifold." I would read only the first paragraph.

(Laughter)

Since then I found another paper and, in fact, this paper appeared after I had done this search, a paper by Marshall Luban and Jim Luscombe. This paper gives you a beautiful explanation of how you get Anderson localization. You should write down that reference.

I can give you some idea of what the problem is. Furstenberg's theorem -- the actual theorem itself, the statement of the theorem, is not too hard to understand -- says that if you have a disordered system, as you go to plus-and-minus infinity, the eigenfunction decays exponentially to zero with probability one. I can understand all those words.

But as an experimentalist, I have problems. First of all, I do not have any infinite-sized systems. Even if you look in the McMaster-Carr catalog, there are not any infinite-sized systems

there. Even if I had an infinite-sized system, what I have to do is march off to infinity and look for something happening with probability one. What does that mean? I do not know.

Anyhow, this paper by Luban and Luscombe explains how you can understand Anderson localization in a finite-sized system with the probability aspect of it built in. It is a beautiful paper. Maybe, in an evening discussion session, I can go over that explanation.

It is a very difficult problem and it is based on the use of statistics of disorder and their effect on eigenvalues and eigenfunctions. In statistical physics, one of the most important parameters is the dimensionality of the system. For this Anderson localization problem, it turns out that in one dimension any amount of disorder will give you some degree of Anderson localization.

The thing is, if you have only a small amount of disorder -- I mean, every real system is going to have some disorder in it -- then this Anderson localization length will be large. If you have a finite-sized system, the Anderson localization length may be larger than the size of your system and, in that case, you will not really see the Anderson localization.

But if you increase the amount of disorder, then the localization length may become smaller than the size of your system and then you will see the Anderson localization.

In three dimensions, it is believed that you have to exceed a certain amount of critical disorder before you see Anderson localization; that is, if you have a small amount of disorder in three dimensions, then the eigenfunctions will still be extended.

If you increase the amount of disorder and you exceed a certain critical amount, then all of the eigenfunctions become exponentially localized; in fact, that is called the Anderson transition.

Two dimensions are the so-called critical dimension. It is difficult to do the calculations in the critical dimension. For a long time the theoreticians were not sure how two dimensions behaved with disorder. They were not sure if there was a finite amount of critical disorder that you had to have in order to see localized eigenfunctions in two dimensions.

We thought maybe we could do an experiment in two dimensions and look for two-dimensional Anderson localization to see if there was an Anderson transition in two dimensions. That was a serious problem in condensed matter physics.

(Transparency)

We took a large aluminum plate, maybe an 8th of an inch thick, a couple of meters or so on the sides, and drive it with a shaker and incite bending waves in that plate. That was our two-dimensional wave medium.

For a potential field we thought we could take a rib structure and attach it to this plate. For a periodic system we thought we could make a rib structure that was a regular array of hexagons, and attach it or machine it into the plate.

Since this is periodic, we could drive this system, look at its normal modes of vibration, and they should be extended Bloch eigenfunctions.

DR. GARRETT: There are flexural waves in the plate and that is a fourth-order equation. That is not isomorphic to Schrödinger equation.

DR. MAYNARD: Right. We also thought about a membrane, or something like that.

DR. GARRETT: It did not bother you that it was fourth order?

DR. MAYNARD: It does not really matter, because the Anderson localization applies to the disordered potential field and it works for more than just the wave equation. If you have a higher order equation, it will still give you disorder.

DR. GARRETT: That was not a problem?

DR. MAYNARD: No. I will explain, we never got to this experiment for a number of reasons. Too much other physics kept popping up right in our faces when we were doing this.

The idea was to have some sort of two-dimensional vibrating structure with a periodic potential field. In this case, with this rib structure, you can take the centers of these hexagons and displace them by random amounts and then you can do a Wigner site-cell construction.

What you do there is you connect the centers and then have a perpendicular bisector of those lines and that gives you an array of disordered polygons. A rib structure like this is a disordered potential field. Now, if we drive this plate and look at its normal modes of vibration, we should see two-dimensional Anderson localization.

In fact, one of the things about Anderson localization is that different eigenfunctions get localized at different sites. In this case, different frequencies, if you put in noise, different frequencies of that noise, different frequency components, get localized in different sites, so we are thinking that maybe you could use a disordered rib structure on a plate to actually filter the noise into narrow band components and then you could use some sort of narrow band noise cancellation technique to reduce the noise; i.e., you use active cancellation in narrow frequency bands at the places where those frequencies get localized.

That is actually still a viable idea. But before we could do this two-dimensional experiment, we thought we ought to do a one-dimensional experiment just to educate ourselves as to how this Anderson localization works. This was before we found the paper by Marshall Luban, so this business of having disorder not giving you disorder in just the wave function but giving this exponential localization, that was a hard pill to swallow.

So we thought we would do a one-dimensional experiment to try to actually see it, literally see it happening.

(Transparency)

We went to the local music store and bought a spool of music wire. This is steel wire, as you use on a guitar. Actually, we found out that harpsichord wire works best, has the highest Q. We took the steel wire, about 50 feet of it, and hung it down the stairwell in the physics building, put a weight at the bottom, so we have a wire with tension in it, and then we put a shaker next to

the wire so that we could excite transverse waves along the wire. That is our one-dimensional wave medium.

For the potential field, for the scatterers, we took little lead masses -- in fact, these are split shot used for fishing lines. We bought these at the local sporting goods store. You can put these little lead masses along the wire and those act like delta function scatters.

The system, the wire with the little masses on it, is a very accurate analog of a one-dimensional wave equation with a Kronig-Penney potential field. This is something you study in solid state and quantum mechanical classes. Here is a nice physical realization of that system.

We could put the masses on here periodically and study extended Bloch eigenfunctions or we could put the masses on here randomly and look for one-dimensional Anderson localization.

To make measurements, there was an aluminum beam running parallel to the wire that had a trolley on it with a C-shaped magnet. That magnet put a field locally across the wire so that when the wire vibrated in that magnetic field, there was a current induced in the wire, so we had one end grounded, and the other end went to an op amp.

With this system we measured the current. You can measure the amplitude and the phase of the vibration of the wire at the position of the magnet. Then you can just translate the magnet down the wire and hook it to a plotter, and you can plot amplitude and phase as a function of position, which is the wave function, the eigenfunction, for the system.

That is something that is just impossible to do in a quantum mechanical system.

DR. SWIFT: Is it important that you chose an anechoic termination?

DR. MAYNARD: No, but we were not sure. This was before we understood Anderson localization. In fact, when you get to the end of that wire with, say, a periodic sequence of masses, there is such a big impedance change from a periodic system to just a bare wire that everything gets reflected anyway. What you do on that end really does not matter.

I think if you want to do calculations, you should try to have something definite on the end of the wire. We could have just clamped it and it would have worked just as well. But we were not sure how that would have affected Anderson localization.

So we could measure eigenfunctions, amplitude, and phase. Before you measure eigenfunctions, you need to know what the eigenfrequencies are. To measure the eigenfrequencies, we would put the magnet in one position away from the driver and we would drive it and sweep the frequency. When the frequency coincides with an eigenfrequency, you get a resonance in response to this system.

(Transparency)

Before I show you the actual data, I want to show you what you would expect to see in a periodic system. In solid state physics, for a periodic system, there are two approaches. There is the free electron approach and the tight binding approach.

It turns out, for our system, the more useful approach is the tight binding approach. Here is the wire -- I have drawn it horizontally in this case, just for the figure -- and you first imagine that the masses are infinite. If the masses are infinite, then it is as though this wire were clamped at the position of the masses.

Now it is easy to understand what the normal modes and the natural frequencies are. They just correspond to fitting an integral number of half-wavelengths between these clamps, so you can fit one half-wavelength, two half-wavelengths. Those are the normal modes.

The eigenfrequencies correspond to an harmonic sequence. The lowest one is the frequency for one half-wavelength and then it is twice that, three times that, and so on, just like the modes of a guitar string. That is if the masses are infinite.

With the masses infinite, it is as though each section of wire here acts as an isolated local oscillator, and this local oscillator has sharp eigenfrequencies.

In the actual system the masses are not infinite; the masses are finite. So when one section of wire moves, the little masses can move a little bit and so the motion of the wire here makes the mass move, which makes the next wire move, and so these sections of wire get coupled together. In other words, the local oscillators get coupled together.

If you have a system of local oscillators that are coupled together, then the sharp eigenfrequencies of those local oscillators broaden out into bands. This is a nice way of demonstrating how you get band structure in solid state physics. In fact, we have a little demonstration that is about so long and has about eight masses on it and you can take it into the classroom and show band structure for a periodic system.

In solid state these are called energy bands and gaps. In acoustics or microwave systems they are called pass bands and stop bands. That gives you this band structure. This happens for the periodic system.

Our system is finite in size. The wire actually had about 50 masses on it. What that means is that each one of these bands has 50 discrete eigenvalues associated with it.

(Transparency)

Now you can see what happens when you sweep the frequency, when you drive this thing and sweep the frequency. You may start out with the frequency in one of these gaps, so you do not get much response -- you actually excite an evanescent wave, so if your receiver is far enough away, you do not really see anything.

So you do not get any response in the gap, but then when you go into a band, you hit these resonances and you get 50 resonances in that band, you go into the next gap and you do not get any response, then you go into the next band and you get 50 resonances, and so on. So now I can show you the data.

(Transparency)

This is for the periodic system. This is the response far from the driver as a function of frequency. Here you drive it and you do not get any response at these frequencies (this is in a gap), but then when you go into a band, you get 50 resonances and then into the next gap. As you go up higher in frequency, you will go into the next band.

This particular band here actually corresponds to fitting approximately one half-wavelength between the masses. The next band would correspond to fitting two half-wavelengths between the masses, and so on. In fact, you can understand what happens in a solid state system now.

If you have an isolated atom, then it has sharp energy levels, but if you take a bunch of atoms and put them together periodically so that they interact, then those sharp levels broaden out into bands that you get in solid state physics. In three dimensions there will be all kinds of structure in that band structure, and group theory is one of the ways you can figure out how these sharp levels evolve into the bands. That is one of the things in Tinkum's book that you can learn about.

In this case, in our simple one-dimensional case, it is easy to associate each band with fitting one, two, or three half-wavelengths between the masses.

You might wonder about what this extra structure is in here, but maybe I will just leave that as an exercise for you to think about.

So these resonant frequencies tell you what the eigenfrequencies are, and now what you can do is you can sit at one of the eigenfrequencies and translate the magnet and get a plot of the eigenfunction for that eigenfrequency.

(Transparency)

That is shown here. These are two different eigenfunctions for two different eigenvalues. This is amplitude as a function of position along the wire. The masses on the wire are actually pretty strong perturbations. The masses do not like to move very much compared to the wire, so they almost look like nodes in the wave function. These points down here are the actual positions of the masses.

If you count, there are about 50 of those points in there. You can also see that between the masses there is one half-wavelength. But there is structure in here, because those local oscillators are now coupled together and the eigenfunctions are not at all trivial any more.

Here, this one has some sort of structure to it that is kind of interesting. If you look in solid state textbooks, you will not see any Bloch wave functions that look like that. For their examples, they always seem to pick points that are at symmetry points in the Brillouin zone. We found that these Bloch eigenfunctions can be quite complicated.

One thing you will notice. In the experiment we do see extended eigenfunctions: This thing has the same nominal amplitude throughout our whole system.

Here it is a little bigger. The problem there is that we use commercial split shot for our lead masses and the masses vary by about 13 percent, so there was some disorder in our system. It

turns out that variations in the mass have a smaller effect than the spacing between the masses. In fact, the spacing between the masses determines the frequency of that local oscillator.

When we were building this thing, we actually had to use a pair of calipers to very precisely set the spacing between the masses. When we did that, we did get a periodic system that has extended amplitude.

(Transparency)

This is really pretty amazing. As I said, those masses are very strong perturbations on the wire. We took the wire and put one mass on it and measured the reflection coefficient and we got almost 100 percent reflection from one of those masses.

But if you put the masses on there periodically, you get the wave going through the whole system and that is Bloch's theorem. It is really pretty amazing. In many ways, the Bloch extended eigenfunctions are more amazing than the Anderson localization. That is addressed to some degree in this paper here, why the Bloch effect is so amazing. That is for the periodic system, we saw the extended eigenfunctions.

That is a good test when you are building one of these systems. You have to be able to build it periodically and see the Bloch eigenstates; otherwise, something is wrong.

Now we wanted to do the disordered system. We took the positions of those masses and changed their positions by just a random amount but within 2 percent. We changed it by random displacement but by no more than 2 percent of the lattice constant. It is actually a very small perturbation.

When we did that, we got a dramatic change in the eigenvalue spectrum. Again, this is response as a function of frequency, but now, instead of getting this nice, sort of regular spacing of the eigenfrequencies, the eigenfrequencies clump up, an irregular pattern in the eigenfrequencies.

One thing that is interesting is you get an eigenfrequency here that corresponds to the gap in the periodic system, so you get states in the gap (this is, of course, what you hear people in solid state physics talking about all the time).

There are also some interesting theorems that treat the nature of the eigenvalue spectrum when you have disorder. One theorem says that if you have an eigenvalue off by itself, that eigenfunction will be highly localized. On the other hand, if you have an eigenvalue that has neighboring eigenvalues, then that eigenfunction will be less localized. (In fact, we saw that in our experiment.)

This is for 2 percent disorder. Now we will sit at this eigenfrequency, translate the magnet, and look at the eigenfunction.

(Transparency)

These five plots are eigenfunctions for the disordered system. You can see we got textbook Anderson localization here. This one is, in fact, that state that appeared in the gap, so that fits that theorem.

The thing to notice is that different eigenfunctions are localized at different sites. Also notice that different eigenfunctions also have different Anderson localization lengths, so the Anderson localization length varies from one eigenfunction to the next, and that corresponds to the universal conductance fluctuations.

In one of those electronic systems, one of those mesoscopic wires, the conductance from one end to the other depends on how much of the amplitude of the wave function you have at the ends of the system. If you have a longer localization length, you will get more conduction. As that localization length fluctuates, then the conductance fluctuates. These would correspond to universal conductance fluctuations.

So we saw how this Anderson localization worked, that, in fact, it did work, which was the point of the experiment, just to educate ourselves that this Anderson localization really happens in a balls-and-springs disordered system.

After doing the experiment, we learned that it was not too hard to do computer simulations for a one-dimensional system that just had the static disorder in it, so we did not really plan on publishing this. But we found out that with this system we could do some more serious physics, and that went like this.

We said suppose you drive the system at an eigenfrequency of, say, 800 Hz, and you get the acoustic energy localized here. Then suppose, as another eigenfrequency, 810 Hz, that has the energy localized over here. We drive this thing at 800 Hz and get the acoustic energy localized here, and then modulate the tension in the wire at 10 Hz, the difference frequency.

When you modulate the tension in the wire like this, the little masses are going like that. That looks just like a longitudinal phonon going through our system. Remember, in our system the transverse waves in the wire are the analogs of the Schrödinger waves for an electron. When you modulate the tension so that the masses move like this -- remember, the masses on the wire correspond to the ions in the metallic crystal -- it is just like having a longitudinal phonon in the system.

With this experiment we could simulate electron-phonon scattering and how it affected the Anderson localization. That is a topic that people were worried about in condensed matter physics.

In the experiment, what we do is we drive it at 800 Hz and get the energy localized here and then start modulating the tension. What you see is that the energy goes back and forth between the two localization sites. One way you can think of it is that the phonon couples the two eigenstates and they beat against one another at the difference frequency of 10 Hz.

MS. GRANT: (Inaudible.)

DR. MAYNARD: No, for low modulation strength it is pretty much just a superposition of the two linear eigenstates.

It is just like two coupled pendula. You can have two coupled pendula with a normal mode like this, or a normal mode like this. If you start it off with just one of them going, that one will slow down, it will pick up over here, and it will go back and forth at the difference frequency between -- the difference in the natural frequencies of those two normal modes.

The same sort of thing is happening here. But now, in the solid state system, you have, say, the electron localized here, a phonon comes by, and the two states are going back and forth like this. But then at some point in time the phonon goes running off to the thermal bath. In fact, that is what the thermal bath is all about; it sends in phonons, it takes out phonons.

Depending on when this phonon leaves, there is some probability that the energy gets left in the second site. This is phonon-assisted hopping from one localization site to another. This was a hot topic on condensed matter physics at the time.

With this experiment, we could simulate it in gory detail. We could measure the two eigenfunctions, amplitude and phase. We could turn on the modulation to any strength. We could get all sorts of high-amplitude things happening in the system.

This last plot down here is a time average of the thing going back and forth like this. From the relative amplitudes of the two states you can determine the hopping probability and you can test theories that use the two initial states, or the initial state and the final state, overlap integrals, and so on, ways of calculating the hopping probability, and we can test it in gory detail.

This is what we published in the paper about the phonon-assisted hopping. It turns out that in this paper this was the first direct observation of Anderson localization in this wire with the masses on it. This was published in *Physics Review Letters*. The experiment cost \$40.00.

So you can do interesting contemporary physics with just balls and string with these analog acoustic systems.

MR. LANDSBERGER: Does it matter what your drive frequency is when you do this experiment? Did you drive it at 800 or 810?

DR. MAYNARD: We drove the transverse waves at the one frequency and then modulated at the difference frequency and the second state picked up. It would be more like a solid state system with electrons if we would just drive it at 800 Hz and get it going in the normal mode that had the electron localized here and then turned on the phonon. Then it would do something and wind up over here.

But there is damping in the system, so you have only so much time to do that experiment, so we drove it just to overcome the damping. The utility of the experiment is still there, because you can still measure the individual eigenfunctions, their overlap integrals, and the hopping probability, even though you are driving it.

DR. BROWN: The randomization is in the mass amplitude, not the base periodicity, correct?

DR. MAYNARD: The important disorder is the spacing between the masses, because that determines the frequency of that local oscillator. If the masses vary a little bit, that does not really change that local oscillator very much. It turns out that when you do the analysis of this thing, which is in this paper, you find that you are dealing with a 2×2 matrix.

The wave equation is second order, so you have a linear combination of two linearly independent solutions. When you look at conditions at one mass, there are two conditions. So to go from one section of the wire to the next, you are multiplying by a 2×2 matrix.

It turns out that the diagonal parts of that matrix have to do mostly with the frequency of the local oscillator. The off-diagonal elements depend mostly on the coupling.

In this Anderson localization business there is a thing called diagonal disorder and off-diagonal disorder. It just turns out that the off-diagonal disorder is much less important than the diagonal disorder, so this is just an example of that in solid state.

DR. MARSTON: With the exception of the lower part of the hopping problem, your measurements deal with steady state phenomena. Is there a simple way to understand, say, in graph C, where I think you were in the gap (the eigenfrequency was in the gap), if you turn on a sine wave drive at time T equal to zero, how long it takes to establish that Anderson localized state?

DR. MAYNARD: Another way of thinking about that is if, instead of just driving it cosine ωT , you send in a pulse or a pulse train? That is actually in here, and I should have made a viewgraph of this.

If you look at Figure 8, this is the periodic system. You send in a pulse (it is an extended pulse, a pulse train). One way to think about what happens when you send in a pulse is you have this periodic sequence of oscillators. A pulse comes in. It pretty much just excites the first oscillator. That one, when it is moving, can transfer energy to the next one, and so on.

In the disordered system, that energy goes down and gets left behind. In fact, in Figure 9 you can see what happens to a pulse going through a disordered system. It just leaves its energy behind and gets smaller and smaller. Some of that energy will build up a little bit more at a localization site, because, after all, the pulse is just a superposition of the eigenfunctions.

You can actually see that in a computer simulation, a pulse going through. That is not too hard to understand. You are exciting one oscillator and then the next and the ones that are localization sites build up a little more. The rate at which it does that is probably difficult to say, but in the periodic system, there you can actually see interesting things happening.

In the periodic system, when the energy goes in and excites this guy, he excites the next one, and so on, this first one transfers all of its energy over to the next one, and the pulse just goes right through, just like a pulse on a string, even though there is this periodic array of scatterers. That is really amazing, if you think about it.

Why should these first oscillators give up all of their energy so that the pulse goes through the system unchanged? That is just really amazing.

That rate you can calculate and it is the group velocity. That is all covered in the gory details in the paper. That is a good point, though.

STUDENT: When the phonon comes through, it can move the electron to another state? I was wondering, with this, when you vary the tension, you get it oscillating between the two, then if you stop it, does it go back to its original?

DR. MAYNARD: Remember, this is the driven case.

STUDENT: You are modulating the tension by 10 Hz, that comes in. So then it is just between 800 Hz and 810 Hz

DR. MAYNARD: No, the frequency does not shift. The position shifts. Yes, because of the damping, if we turn off the modulation and we are driving it at 800 Hz, it will just go back to the 800 Hz, but that will be at the decay rate that is determined by the damping of the system.

STUDENT: Is that the same observation in the phonon-assisted hopping?

DR. MAYNARD: No, because -- let us see, in that system I guess there is damping, also. The thing is, there, if you try to go to low temperatures so that the bigger effect is the hopping, the electron really does not lose its amplitude to some other form of energy. There the existence of the electron does not decay. Once it hops to this site, it will just stay there and resonate at that site.

Okay, so the experiment that we originally wanted to do was two-dimensional Anderson localization, but we found that there was a lot of extra physics in the one-dimensional experiment, so we spent some time on that.

And even then, before we could get to the two-dimensional Anderson localization, we got derailed, again, because quasi-crystals were discovered.

(Transparency)

It used to be thought, in solid state physics, that there were two basic types of solids: crystalline and amorphous. In an amorphous material the atoms were more or less randomly positioned and had isotropic, homogeneous properties. A crystal is quite different.

To make a crystal, you take a unit cell and you repeat it periodically to fill all space. In order to fill all space without having any gaps, only certain shapes are allowed. In fact, there are 14 different shapes and those form the basis for the Bravais lattices.

Because there is only a finite number of shapes that are allowed for filling all space, only certain rotational symmetries are allowed. In fact, fivefold rotational symmetry is not allowed.

Here is a quote from Kittel's solid state physics book. It says: "A five-fold axis of symmetry cannot exist in a lattice because it is not possible to fill all space with a connected array of pentagons."

Some time ago, Schechtman, at what was then the National Bureau of Standards, made an aluminum alloy, got some small crystallites and put them in an x-ray diffraction machine, and he got this diffraction pattern. He got 10 spots in a circle, which means he had a mirror plane and fivefold rotational symmetry. He showed this to people and, of course, they all laughed at him. "Ha, you cannot have fivefold rotational symmetry. You've done something wrong."

So Schechtman and other people re-did the experiment, made better quality samples, and measured the x-ray diffraction, and the spots just kept getting sharper and sharper.

A sharp spot in a diffraction pattern means that you have long-range order. So here Schechtman found this material that had long-range order and fivefold rotational symmetry. It turns out that mathematicians and physicists knew that this could happen. It is impossible to have fivefold rotational symmetry and long-range periodic order, but you can have fivefold rotational symmetry and another type of long-range order, which is called quasi-periodic.

(Transparency)

Here are some pictures of these quasi-crystals. These are maybe a couple hundred microns in size and you can see there are facets. This is a macroscopic crystal. It has got facets in a fivefold pattern, so it has global fivefold rotational symmetry. This is this aluminum alloy.

(Transparency)

Here is another example. You can see the fivefold symmetry here. This one was made by the Alcoa Company and apparently they make it especially for the Chrysler Corporation.

(Laughter)

(Transparency)

In one dimension it is not too difficult to understand quasi-periodicity or quasi-crystals. Let me give you an example. Suppose I take a line and I put dots on it periodically with a lattice constant a . If I did a diffraction experiment on this system, in other words, took its Fourier transform, I would get a sharp line at π/a . I get a sharp line indicating that I have long-range order (it is long-range periodic order, in this case).

Another way of looking at it is, over here, I have only one characteristic length and so this line is at π over that characteristic length.

On the other hand, I could take a line and put dots on it randomly. If I take the Fourier transform of that, I get just a broad distribution in the spectrum. That is because in the random array of dots there is no one characteristic length.

But now suppose I take a second line and put dots on it periodically with a lattice constant b , so its Fourier transform is a line at π/b . Then suppose I superimpose these two lines. I will get a line with dots on it that may not look all that different from the line with the random dots on it.

If these two lattice constants, a and b , are commensurate, that is, if their ratio is equal to a rational number, which is equal to a ratio of two integers, then it turns out that there will be some pattern of dots here that will repeat periodically.

If the ratio of integers is a ratio of two large integers, then that unit cell here, that pattern, may be very large, but eventually it will repeat periodically and, in that case, you can use Bloch's theorem and other powerful group theory theorems for periodic systems. That is if the two lattice constants are commensurate.

On the other hand, if the two lattice constants, A and b , are incommensurate, that is, their ratio is equal to an irrational number, then there will be no pattern here which repeats periodically because of that irrational number.

But if you take the Fourier transform, you get two sharp lines, because this was just a linear superposition of these two periodic structures. You get two sharp lines and that indicates you have long-range order, but it is not periodic order. It is not periodic because these spatial frequencies are not commensurate.

That is the notion of quasi-periodicity in one dimension.

Quasi-crystals in higher dimensions are a bit more sophisticated than that. To explain more about quasi-crystallines in higher dimensions I am going to show another way of making a quasi-crystalline sequence in one dimension.

In this second method you start off with a periodic array in a higher dimension; in this case, a two-dimensional square array of dots. Then you intersect that higher dimensional periodic structure with a lower dimensional surface. In this case, I intersect the two-dimensional array with a one-dimensional line.

Then you define a window and within that window you project the points down onto the line. If this lower dimensional surface has direction cosines that are irrational, then the sequence of dots on the lower dimensional surface will form a quasi-crystalline sequence.

The thing has long-range order because it has been projected from a periodic higher dimensional structure -- that gives it long-range order. But because of those irrational direction cosines, there is no pattern that repeats periodically.

In this one-dimensional example there is a special case where, if the tangent of this angle is equal to the Golden Mean, which is the square root of 5 plus 1 over 2, then the sequence of points that you get on that line corresponds to a sequence of Fibonacci numbers.

You probably know from high school that Fibonacci numbers have all kinds of miraculous properties. By the same token, quasi-crystal formed by having this irrational angle of the Golden mean also has pretty amazing properties and actually allows you to prove rigorous theorems and so on.

The Golden Mean, the square root of 5 plus 1 over 2, is also known as the Devine Ratio and it is also known as the most irrational number. I think that is because if you do a continued fraction approximation of the Golden Mean, it is the continued fraction that converges the slowest. Anyhow, you get miraculous properties for this particular case.

In higher dimensions, in order to make a three-dimensional quasi-crystal, you have to start off with a six-dimensional periodic lattice and intersect it with a three-dimensional surface. I am sure you will have no trouble at all imagining that. When you do that, you have a window, project points onto that three-dimensional surface, and you get a three-dimensional quasi-crystal.

For two dimensions, you start off with a five-dimensional periodic structure, intersect it with a plane with irrational directional cosines, and project the points, and you will get a pattern that is popularly known as a Penrose tile.

(Transparency)

These structures were studied by Penrose a lot and they were written up in *Scientific American* in Martin Gardner's column. They are quite popular. Also, *American Scientist* has a feature article that is very good. I have been told that there is even a company in Texas that makes bathroom tiles that you can use to tile your floor with one of these Penrose tile patterns.

This is a standard one. It has fivefold rotational symmetry. These tiles fill an infinite plane without leaving any gaps.

To make an ordinary crystal, you are allowed to use one unit cell and repeat it to fill all space. To make a quasi-crystal, you are allowed to use more than one unit cell. In this particular pattern there are two unit cells: there is a skinny rhombus and a fat rhombus. Those will fit together to tile the plane and wind up having fivefold rotational symmetry and no periodic order.

In order to tile the plane, to fill the plane without leaving any gaps, you have to follow some very definite rules. If you make any mistakes, you are likely to get to a point where you cannot fit a tile in. Because you have to follow these very definite rules, these structures have long-range order.

But the ratio of the areas of these two tiles in this case is the Golden Mean, so there is an irrational number involved in the length. Because of that irrational number, there will be no pattern in this structure that repeats periodically.

Even though there is no periodic structure and no Bloch's theorem, there are some very amazing theorems associated with these structures. If you take these tiles and draw lines on them in some way (I am not sure how you do it) and then put the tiles down according to the rules, you will find that the lines that you drew on the tiles also form a Penrose tile pattern, but they will be scaled down. That is called inflation or deflation symmetry.

There is another theorem that says pick any pattern in here with some nominal diameter, then within two diameters, and usually within one diameter, you will find that pattern again. In this

particular one I actually have a global fivefold symmetry point, so it is easy to find that the pattern repeats.

But if you are way off the symmetry point and you find any pattern, within one or two diameters you will find that pattern, again. That is pretty amazing. I cannot even imagine how they discovered that, let alone proved the theorem. That is called Conway's theorem.

I think that theorem has very important consequences for the properties of these quasi-crystals.

MR. PRATHER: You said there is only one?

DR. MAYNARD: Yes. This one actually has a fivefold symmetry point. Nothing repeats periodically here, so if you are way off away from that point -- I am not even sure you have to have a global fivefold symmetry point. It is only when you take its Fourier transform that you will see the fivefold symmetry.

(Transparency)

There are three fundamental types of solids: periodic structures, or crystals; random structures, or disordered solids; and now, quasi-crystalline solids. Now there is a fundamental question. Suppose you have a wave equation with a potential field that has a quasi-crystalline symmetry? How does that quasi-crystalline symmetry affect the eigenvalues and the eigenfunctions? Very basic question.

(Transparency)

It turns out that in one dimension you can prove rigorous theorems. That is because of those Fibonacci numbers. In one dimension you can show that the eigenvalues form a cantor set, and that undoubtedly has something to do with the fact that M is a manifold, but anyhow, you get a cantor set.

(Laughter)

You get bands and gaps everywhere and the ratio of the amounts of bands and gaps is the Golden Mean.

The eigenfunctions for a one-dimensional quasi-crystal can be either extended, localized, or so-called critical eigenfunctions, depending on parameters in the system. That is for one dimension. You can prove rigorous theorems.

(Transparency)

For two dimensions and three dimensions there are no rigorous theorems. There is no quasi-Bloch's theorem. The problem is that the system has long-range order, but it is not periodic, so you cannot use Bloch's theorem, and it is not disordered, so you cannot use the statistical techniques that you use in Anderson localization length. So far, no one has been able to prove any theorems for these higher dimensional quasi-crystals.

Without any theorems to make predictions, we decided we would just do a measurement to see how two-dimensional quasi-crystalline symmetry affects eigenfunctions and eigenvalues.

(Transparency)

For the experiment, once again, we wanted to use the tight binding approach, where you start off with a local oscillator that has nice, sharp, well-defined frequencies and then have a bunch of these coupled together so that instead of a sharp frequency you get a whole spectrum of frequencies. We could look at that spectrum of frequencies and see if there was any structure that might reflect the quasi-crystalline symmetry.

We wanted to put a local oscillator in here in this pattern and then couple these local oscillators together so as to maintain the fivefold or the quasi-crystalline symmetry of the thing.

In these experiments, one of the things that is very important is you want low damping, because, remember, the theorists like to call it long-range phase coherence, so you want to be able to have a wave that can probe this whole structure and have interference effects, and so on, without much damping (if it damps out, you are not going to get much of an interference effect and you are not going to see as much structure). So you need low damping or a high-Q local oscillator.

If you want a high-Q local oscillator, it is hard to beat the old tuning fork. I can start this going and it will ring for quite a while. It will do that even though this is a lossy support where I am holding it. The two tines work against one another, so there is not so much motion here and energy cannot leak out.

Also, you cannot hear it, and that is good; that means it is not losing energy by radiating sound out into the room. This is a nice high-Q local oscillator, 10,000 or maybe approaching 100,000 for the Q.

(Transparency)

After buying a couple hundred feet of music wire, we went back to the music store and bought 300 tuning forks. We took a large aluminum plate, about a meter and a half on a side, and had the shop drill holes in it in a Penrose tile pattern, and epoxied the stems of the tuning forks into the plate in that Penrose tile pattern.

There is a local oscillator with a frequency of 440 Hz, quite accurately set. We have them in the Penrose tile pattern and now we want to couple them together. For coupling, what we did is we took an arc of steel wire and spot-welded the arc of steel wire between the tines of neighboring tuning forks.

When the tines move, that arc of steel wire will spring. Actually, getting the coupling right was really a difficult thing, because you need some coupling system that does not have much damping. At first, we took these arcs of wire and just epoxied them to the tuning forks, but the epoxy was way too lossy. What we did, then, was we spot-welded them and that worked quite well, maintaining the q.

For the pattern of those coupling wires, here is one of those rhombuses -- like this. The tuning fork was epoxied in there so that the two tines were across the short diagonal. Then the four sides of the rhombus identified four nearest neighbors. Each tine, then, had two coupling wires going out to the nearest tine of the neighboring tuning fork.

It is actually a complicated system; it is like a quasi-crystal with a diatomic molecule at the sites. But the fact that this thing has Penrose tile symmetry is inescapable; that is built into the system, guaranteed.

For measuring this we took an electromagnet and put it next to one of the tines and then just put an AC current in the electromagnet to drive the thing. To measure its response, we went back to the music store and bought a half-dozen electric-guitar pickups. We put the electric-guitar pickups at random places in here next to tine and measured their amplitude of response.

In fact, we moved them around and measured 20 different spots to make sure there was a good chance we did not just sit on a node or something. Then we would sweep the frequency and measure the response and get these resonances that would tell us what the eigenfrequencies were.

(Transparency)

Here is a plot that shows the eigenfrequencies. Here is the 440 Hz. It actually falls in a gap. With this sort of diatomic system there is an acoustic mode, where the tines can move together and there is an optic mode, where they can move apart. That gives you two branches in the eigenvalue spectrum.

The whole spectrum is only a little over 100 Hz wide and there are 300 lines in there. In order to resolve all those lines just a fraction of a hertz wide, we really needed that high-Q local oscillator, so the tuning fork was a serious concern here.

Paul Chakin, at Princeton University, had a postdoc working on this problem for a year. He had a rubber membrane for a two-dimensional wave medium and he was putting little epoxy dots on it in a Penrose tile pattern. But that experiment never worked, probably because there was too much damping in it. So thinking of using the tuning fork was a good thing to do.

If you look at this eigenvalue spectrum, the local oscillator, 440 Hz, coupled together in the Penrose tile pattern, you get the spectrum and the spectrum does have structure in it -- there are these gaps. If you measure the width of these bands and gaps and take the ratio of those widths, you get the Golden Mean. The Golden Mean is built into the system, or does show up in the system.

People had done computer simulations of two-dimensional quasi-crystals and they did not see this effect. There is a good lesson here. In those computer simulations they used a hopping Hamiltonian. The real problem is the Schrödinger equation, which is a wave equation. But they used a hopping Hamiltonian, it was just a big matrix, and the matrix element was either one or zero, depending on whether or not there was a nearest neighbor in the Penrose tile pattern.

That has the topology of the Penrose tile but they threw out the wave equation. When you do that, you do not see this effect. In fact, the way you can understand this is that you do not throw out the wave equation.

(Transparency)

Another way to say that is our coupling wires have a finite mass. There is a finite speed of propagation in the coupling wire system. This hopping Hamiltonian you can get if you assume coupling by massless springs, so that the coupling system has an infinite speed of sound.

But if you have a finite speed of sound, then as you change the frequency, as you are sweeping the frequency of the system, the wavelength is changing. As that wavelength changes, it probes all the length scales in your system. Of course, in this system the Golden Mean is in there all over the place, so the Golden Mean shows up in the length scales and, because of the finite speed of sound, it shows up in the frequency.

Because of Conway's theorem, patterns repeat -- here is a pattern and that repeats a large number of times. Suppose you are at a frequency where you can fit an integral number of wavelengths around this circumference here. It is going to like that and lock up. Then as you sweep frequencies, these sorts of structures are going to produce bands and gaps in the eigenvalue spectrum.

In the computer simulations they essentially threw out the wave equation and did not see the structure. You have to remember that the hopping Hamiltonian is a model and the Schrödinger wave equation is a law of physics.

(Transparency)

In addition to measuring eigenvalues, you want to measure eigenfunctions. To measure the eigenfunctions, you want to measure the vibration of all the tines of those tuning forks. We took the tuning fork quasi-crystal and mounted it up on the wall in the physics building and, on the end of a tine, put a little mirror. Then you shine a laser onto that mirror and it reflects and goes across to the other side of the room.

Even when you are driving this thing at resonance, the tines of the tuning forks do not move very much, you cannot see anything (you can sort of feel it). But now, if you shine the laser on it and it gets reflected over, then as the tine moves you have this big optical lever arm. You can actually get a few-centimeter streak of the laser beam on the opposite wall.

We had the system mounted on the wall, we had the little mirrors on the end of each tine. The laser, then, went through a scanner so it would scan over the whole tuning fork quasi-crystal and then the laser beam would hit a mirror, get reflected to the other side, and then we would darken the room, open a camera shutter to a pattern on the opposite wall, and then scan the laser.

The film would pick up the streaks of all the reflections from the tines and from that you could see the amplitude and, actually, the polarization of the vibration of all of the tines.

The problem is that in order for this to work, the little mirror on the tine has to be adjusted so that the laser beam goes to the corresponding position on the other wall. The poor graduate student was up there for months adjusting 300 little mirrors.

(Laughter)

But he had to get it only approximately right. It actually worked.

The problem with that whole measurement, though, is that when you have a system and you may have little perturbations in it, the eigenvalues are not very sensitive to those perturbations; there is this adiabatic theorem that people like Mo Greenspan knew all about.

On the other hand, eigenfunctions are very sensitive to perturbations. When we were measuring the eigenfunctions, those spot welds from those coupling wires did vary by some amount and, as a consequence, the eigenfunctions are not great, but there were some qualitative features.

(Transparency)

We did see some eigenfunctions that had pretty much uniform distribution of the amplitude, some had different patterns, and there were some that had quite localized amplitudes. But the result with the eigenvalues was very definite, finding the Golden Mean in the eigenvalues.

This is a popular experiment, using tuning forks. Again, the whole experiment cost a couple hundred dollars at most in this case.

(Transparency)

It got featured in *The New York Times* science section, which was kind of fun.

DR. MARSTON: How important was it that the tines were in a plane? That is to say, at least from your figures, it appears that when you cemented the tuning forks in the holes that you lined these tines up in a plane. Was that the case or not?

DR. MAYNARD: No, no. I took some poetic license in the figure here. The tines were actually oriented -- if you look at each rhombus here, the rhombi are oriented in some Penrose way -- but the tines were oriented across the short diagonal of every rhombus.

DR. MARSTON: So you always used that rule?

DR. MAYNARD: Yes, a definite rule. Then the four sides identified the four nearest neighbors. There was a very definite pattern in the coupling.

DR. ATCHLEY: What about the length of the wires?

DR. MAYNARD: If you look at the Penrose tile pattern and of all the possible ways that these tiles can go together, the distance between the tiles took on one of, I think it was, four different values. Of course, those lengths had the Golden Mean in them -- I mean, there is a little difference, because there is some separation of the tines and so on. There were only four different lengths.

In fact, to make it uniform, we made a little jig where we could make the four different lengths very uniformly and then spot-welded those in.

MR. LAVRENTYEV: How were the forks excited?

DR. MAYNARD: We used an electromagnet and put an AC current through it. That would pull one of the tines and then the whole thing would resonate.

MR. LANDSBERGER: This is what kind of a surface here?

DR. MAYNARD: It was about a 3/4- or 1-inch-thick aluminum plate and epoxied by the stems.

MR. LANDSBERGER: And you did not excite that plate?

DR. MAYNARD: No. In fact, one of the ways we thought they might couple is through the stem and through the plate, but that was very weak -- very weak.

We tried a bunch of different schemes for coupling them and this was the one -- in fact, we were hard-pressed to get enough coupling to broaden out the spectrum enough to where we could actually identify the lines and gaps and so on.

STUDENT: So what did it sound like?

DR. MAYNARD: It sounded really neat. We had people coming in all the time, photographing it and micing it. You could pluck it and it would give this very eerie sound. It might have been Dr. Hargrove's idea, but somebody suggested that you take your hi-fi, play music into the drive, take it out, amplify it, and put it through the speakers, so filter your music through this spectrum and see what it sounds like. We never did that.

(Transparency)

We have discussed systems that are disordered, but also nonlinear. I have talked about waves in disordered systems and it has been a real hot topic, a real frontier area, in physics, Anderson localization, universal conductance fluctuations, normal electron persistent currents, lots of *Physics Review Letters*, a Nobel Prize, and two Buckley Prizes on waves in disordered systems.

At the same time, but sort of independently, there has been a lot of frontier research on nonlinear systems, solitons and chaos. There have been books written about chaos, a lot of popular books on chaos. Chaos is the source of a lot of humor in physics.

(Transparency)

Somebody sent me this cartoon. It says, "What led you to the mathematics of chaos, Dr. Maynard?" I do not know how they got my name, but it is amazing, it is a pretty good likeness, too.

(Laughter)

Now, with so much activity going on in waves in disordered systems and nonlinear waves, there ought to be a lot of exciting research going on in the overlap of those two fields, where you have a system that is both disordered and nonlinear.

(Transparency)

If you look at the overlap of those two fields, you find that, in fact, there is not much going on. The reason is because this is a very new field and it is also very difficult. The first meeting on disorder and nonlinearity was in 1989 and, since then, I think there has been one other meeting, and maybe another one.

At those meetings, every single paper was a theory paper, so there is lots of opportunity for doing interesting experiments here. We are doing some and there are plenty more that can be done. The experiments are also quite difficult.

(Transparency)

Now, for a system that is both disordered and nonlinear, there is a fundamental question, and that is, does the nonlinearity weaken Anderson localization? If you have disorder, at least in a one-dimensional system, if you have disorder, you will get Anderson localization. But if the waves are nonlinear waves, does that destroy the Anderson localization?

If you look in the literature, there are only about eight groups that have tried to address this question. A lot of the people here are hard-core mathematicians -- again, using manifolds and things. They have tried to answer this question. Some of the theory groups say, yes, nonlinearity weakens localization. Some say no, it does not. Some vote both ways.

(Laughter)

You might think there is some fudging going on here, or some controversy going on here, but it turns out there is no controversy. It is a nonlinear system and nonlinear questions sometimes do not have unique answers. Or, in this case, the answer that you get depends very precisely on how you ask the question.

(Transparency)

In a one-dimensional acoustics system there are two ways you can analyze the system. You can do a normal mode analysis or you can do a pulse analysis. In a normal mode analysis you take the system and you drive it at one end, cosine ωt at some frequency ω . Then you look at the amplitude that gets transmitted and you look at that as a function of the frequency ω .

When the driving frequency corresponds to one of the eigenfrequencies, you will get a resonance in the response to the system and you will get a peak in the transmission. In the normal mode analysis you look at a transmission spectrum.

In a pulse analysis, you take a system and you send in a sharp pulse. That pulse goes into the system, it rattles back and forth among all the scatterers. At the exit point of the system, if you look at the displacement as a function of time, you will get some kind of sequence of pulses

coming out -- maybe even more random, like this. So you are measuring displacement as a function of time at an exit point.

In a linear system, these two things are just Fourier transforms of one another. Whether you do a normal mode analysis or a pulse analysis -- or, I guess, I just thought of this, whether you do RUS or whether you do Moises' pulse measurements -- mathematically, at least, they are equivalent.

But in a nonlinear system that is no longer true.

DR. MIGLIORI: I am sorry, you lost me. You are putting in a pulse --

DR. MAYNARD: A delta function, not a pulse train.

DR. MIGLIORI: A delta function in pressure, the Fourier transform of that would be all frequencies.

DR. MAYNARD: That is right. That is this.

DR. MIGLIORI: Oh, so you are sweeping? I see.

DR. MAYNARD: The whole collection of frequencies, yes.

In a nonlinear system, during a normal mode analysis, you could get a different answer from doing a pulse analysis.

(Transparency)

Let me go to this normal mode analysis and what the theory predicts for the nonlinearity and the disorder. The theory predicts that for a normal mode analysis the nonlinearity does not destroy the Anderson localization. But the theory was not quite complete and there was kind of an open question in the theory.

It turns out that our experiment with the masses on the wire was suited to address that question. Let me go over nonlinearity in a string in a little detail.

Here is a string. It is stretched to some tension

-- say T_0 . You can put a transverse displacement on the string and it satisfies the wave equation. The speed of the wave is the square to the tension over the mass-per-unit length.

But now suppose you give the string a finite transverse displacement. In that case, the arc length of the string has changed. If you stretch the string, then the tension changes (let us say you clamp it at this point). The change in the tension is given by this expression -- it has got the arc length in it -- but, anyway, it has got the amplitude of the displacement in here.

If you have a finite amplitude transverse displacement, then the displacement appears here -- you could do a Taylor series expansion and you will get a term with x_i in it and there is already a factor of x_i here, so you get a nonlinear wave equation. (In fact, you get the nonlinear Schrödinger equation.) This is something that hard-core mathematicians spend their careers studying.

One thing I want to point out here is that if I give this thing a local transverse displacement, then the tension is going to change here, but that change in tension is going to propagate along the

wire, if this is a steel wire, at the speed of sound in the steel wire, which is a lot faster than the speed of propagation of that transverse displacement.

In a steel wire, if you give it a local displacement, the change in tension appears on the whole wire almost instantaneously.

MR. HALLEJ: How do you do the displacement now? Do you tap it just at one spot or is this --

DR. MAYNARD: This is just theoretical, it is hypothetical.

MR. HALLEJ: This displacement is felt across the wire the minute you do it? Are you picking it up and then letting it go, or do you have to hit it?

DR. MAYNARD: You could hit it, give it an impulse. Another way you could think of doing this is like putting in a glass tube and then try to get Keolian to shatter the glass tube, so you get some initial condition with this shape.

Even though the finite displacement is localized here, the change in the tension is over the whole string. If you put that in here, you find that you have a global nonlinearity, where a localized finite displacement gives you a global nonlinear effect. That is a little different from surface waves on water and so on.

Here is the idea. If you put masses on here and then have it disordered, then, at least in the linear system, you will get Anderson localization at some site. Maybe at a different frequency you get Anderson localization at a different site.

Now, suppose you have Anderson localization here and you have finite transverse amplitude. That finite transverse amplitude is going to modulate the tension in the whole wire. When you take a string or a wire and do this to it, just stretch it, it is unstable to the formation of transverse vibrations. If you drive it, as Dr. Keolian showed in his demonstration, at twice the frequency, you can excite the resonance at its frequency for transverse resonance.

What we were hoping would happen here is that Anderson localization at one site would modulate the tension in the wire and that modulated tension would then parametrically excite a different Anderson localization state. So we would have nonlinear assisted hopping from one localization site to another. That, in fact, is the case that could not be treated in the theory paper.

One thing you might think is that, well, you are parametrically exciting it at this frequency, but this eigenfunction localized over here was at a different frequency. Again, as Dr. Keolian showed us -- see, in a linear system you have two different eigenfrequencies for this state and that state, so if this one is running, you might think it would not parametrically excite that one, because it is at a different frequency.

But in the nonlinear system you get the bent tuning curve -- like this -- so at one frequency either one of these states can run. This one can run at this amplitude and this one can run at that amplitude at the same frequency.

We actually looked for that. We tried to find that in the experiment. We looked for two states that were close together -- yes?

DR. KEOLIAN: Also, if it is parametric, there is a finite width at which you can excite a mode.

DR. MAYNARD: Anyhow, we looked for this effect. We tried to find two localized states that were far apart in space but close together in frequency. We would drive one and see if it would parametrically excite the other or have some nonlinear hopping between one localized site and another. It would not do it.

(Transparency)

In fact, here are some data. This, now, is the transmission spectrum for a disordered system, so the eigenvalue is all over the place here. This is at one drive level. Then we increased the drive level to get a finite amplitude displacement and the drive is increased by almost three orders of magnitude here.

Then the curves are normalized by the drive level, so that if this were a strictly linear system, all of these curves would look exactly the same. But they are changing. If you look here, this peak here gets bigger, bigger, gets quite big there, but then it drops off.

This one here, this one starts getting bigger, but then it drops off.

If you integrate under this curve to get the total transmission, then you find if you normalize to this low-amplitude one that as you increase the amplitude and normalize things, you actually get less total transmission.

(Transparency)

If you look at an actual eigenfunction -- this is normalized drive increased by a factor of 50 -- here is a localized state here. As you increase the drive, maybe one of the sections is enhanced by the nonlinearity, but the localization length does not really change very much. In fact, if anything, the states tend to get more localized, so you get less transmission.

Instead of one Anderson localized state parametrically exciting another, it seems to parametrically excite itself, it digs a deeper hole for itself. That sort of thing happens in a disordered system. That is the story for the normal mode analysis; they stay localized (that agrees with the theory).

(Transparency)

Pulses are something else. Let me talk about a pulse in a linear system. There are some people who tried to study pulses in disordered systems and I think they made a big mistake. What they would do in trying to study a pulse is they would send in a pulse and try to keep track of it as it reflected off of the scatterers.

That is hard to do, but they found an easy way to do it.

(Transparency)

They came up with a special system of scatterers, where they would change the density of the medium between the scatterers, so there was a change in the impedance and also a change in the speed of sound, so they get reflections, but then they adjusted the distances between the scatterers so that the transit time across was the same.

If you think about it, that is a medium with interfaces in it, but it has this special feature, that each section has the same transit time. You can make it random. You can send in pulses and because the transit time is equal for each section, it is a lot easier to keep track of the pulse.

When you do that, you get this very interesting pattern of the pulses coming out. What people said is that, "Oh, well, we have a pulse going in and this looks like a broadened pulse coming out." But that is completely erroneous. This is just a complete artifact of this special medium. If you do a random medium and you send in a pulse, you get sort of a mess coming out.

What they were trying to say is that a pulse diffuses through a one-dimensional medium, but that cannot happen. The whole idea of Anderson localization -- the title of Anderson's paper was "The Absence of Diffusion in Certain Random Systems." So when you have disorder in one dimension, you cannot have diffusion; you can have only Anderson localization.

The whole thing with this special medium was a red herring.

(Transparency)

The proper way to do a pulse, at least in a linear system, is to use the Fourier transform of the transmission spectrum. To calculate the transmission spectrum, you start at the exit end, because you know at the exit end you have a wave going only that way, then you work your way back up the system and, as you go across each scatterer, you are multiplying by a 2×2 matrix until you get to the other end. So you have a product of 2×2 matrices.

If this is a disordered system, you have a product of random matrices. The title of Furstenberg's paper was "The Product of Random Matrices" (that is what the theorem was about).

So if you have a sufficiently long system and a product of random matrices, then by Furstenberg's theorem you are going to get Anderson localization.

For a pulse, even though you use this to calculate the transmission spectrum, this is just a Fourier transform of that, so if Furstenberg's theorem guarantees you get localization in the eigenfunctions that go into this, then if you follow the pulse, what is going to happen is you are going to see this pulse, the amplitude, exponentially decay (actually, its energy gets spread out as it goes along). So a pulse will exponentially decay as you go through this linear system.

Even though the pulse is localized, this proof with Furstenberg's theorem requires that you actually satisfy conditions over the whole system, because to get this you have to multiply by these 2×2 matrices from one end to the other. So even for a pulse you really have to satisfy conditions over the whole system to satisfy the wave equation.

Now, for a nonlinear pulse, the thing about a nonlinear pulse is that it has extra degrees of freedom. In fact, for a soliton, to treat a soliton mathematically, you use inverse scattering theory. In an inverse scattering problem you have eigenstates and eigenvalues.

To make up a soliton, it is somehow related to some combination of eigenstates from the inverse scattering problem. For a soliton, the extra degree of freedom would be the number of eigenstates related from the inverse scattering problem. There are mathematical definitions for these things.

A nonlinear pulse has this extra degree of freedom, so now in your system, with a nonlinear pulse, you can adjust these extra degrees of freedom so that you only have to satisfy conditions locally. In other words, a nonlinear pulse has its own characteristic length. For a soliton, in fact, it is the width of the soliton (another thing you can quantify mathematically).

Now we have two lengths in the problem. For a system with a nonlinear pulse and disorder there are two lengths: There is the nonlinear length and there is the Anderson localization length.

Now, with a pulse, interesting things can happen, depending on the relative size of those two lengths. This is what the theory predicts for a pulse. This is the natural log of the transmission as a function of the distance that the pulse has traveled.

If you have a strong soliton, or a strong nonlinear pulse, so that the nonlinear length is less than the Anderson localization length, then that pulse does not probe enough of the disorder to satisfy Furstenberg's theorem and so it does not know about Anderson localization. For a short nonlinear pulse there is no Anderson localization and so a strong nonlinear pulse will go right through a disordered system with no decay.

On the other hand, if you have a weak nonlinear pulse, such that the nonlinear length is bigger than the Anderson localization length, then it learns about the Anderson localization and a weak nonlinear pulse will exponentially decay as it goes through a disordered system.

But then there is an interesting case where the two lengths are comparable. In that case the theoreticians predict that this intermediate nonlinear pulse will go along, decaying a little bit for a while, but then there will be a break and it will decay exponentially.

We wanted to do an experiment to try to observe that effect.

(Transparency)

For this kind of nonlinearity for a pulse we needed a local nonlinearity, not this global one for the string. For local nonlinear pulses, solitons, it is nice to use surface waves on a fluid. In a surface wave, the speed of the wave depends on depth. But when the surface wave is present, it modifies the depth -- like this. If you put that in the wave equation, you naturally get a nonlinear wave equation.

In addition to this nonlinear pulse we wanted to have some Anderson localization. In other words, for the weak or linear pulses we wanted to see the Anderson localization. That means you need long-range phase coherence, so you want low damping, again.

For our surface wave we used a surface wave on superfluid helium. This is actually a thin film of superfluid helium; it is so thin it is only six atomic layers thick, or about 20 Å thick. For such a thin layer, normal fluid is just locked to that substrate and cannot move at all. But the superfluid, which has no viscosity, can still move to make a wave. That type of wave, in superfluid helium, in a film of helium, is called third sound. There are some complications with third sound, but I need not go into those here.

The other thing about third sound is that the restoring force is not gravity, as it is for a shallow water force; the restoring force is a van der Waals force. But anyway, there are well-established techniques for doing third sound.

In the experiment you need a substrate for holding the superfluid film -- we used a microscope slide. For a transducer we use a thin resistive strip here and put a pulse of current through it so that you actually heat the film and that causes a little bit of it to evaporate and thins the film. Actually, instead of having a pulse that mounts up, you have one that goes down. There is a temperature component associated with the pulse and so you can detect the pulse with a thermometer. These are third sound transducers, line sources, and we try to drive plane waves down this substrate.

Then we need an array of scatterers. For the scatterers we took a diamond wire saw and cut a notch in the microscope slide. The biggest effect is that it produces a sharp corner here. This is much bigger than the thickness of the film here. When the third sound has to turn that sharp corner, there is essentially an impedance change for third sound and you get reflected third sound. These are scatterers for third sound.

DR. GARRETT: Are you sure that when you heat that resistance wire that the films thins by evaporation? Doesn't it thicken by a fountain effect if the superfluid comes rushing in?

DR. MAYNARD: I think the evaporation is the bigger effect. I am not positive. Bob Halleck did a measurement and that was the impression I got.

DR. GARRETT: I thought it mounted up where you get it hot, because the superfluid comes and brings -- but maybe there is a critical thickness.

DR. MAYNARD: Could be.

MR. HALLEJ: If you put helium in a cylinder and put a wire on the top, then two waves will go around the sides.

DR. MAYNARD: Right, that was one of Rudnick's old experiments. But there was another experiment where they used these transducers but they also measured the film thickness directly,

and so they looked at a correlation between the temperature swings and the actual displacement of the film. Somehow I thought it was the other way around, but I could just have it backwards.

(Transparency)

Before we did the experiment with the scatterers in there, we wanted to look at the nonlinear properties of a third sound pulse by itself. This is now a bare substrate with no scatterers. I am looking at third sound as a function of the amplitude of the drive.

The thing is driven over here. Then there is a time of flight down the microscope slide and then there is a received signal at the receiver. As you increase the drive level, the thing just scales up.

(Transparency)

If you go to a large enough drive level, what happens is this pulse -- here -- stops growing. In the third sound there is the elliptical motion of the fluid. The velocity of the fluid in that elliptical motion exceeds the critical velocity and then it cannot go any faster, but there are still other things that can happen.

What happens here is that after this thing saturates, you keep increasing the drive, a second little feature grows out of this, moves along as you increase the drive, now, and you can see it easily if you take two curves at a slightly higher drive level -- so the thing has moved over -- and when you subtract those two you can actually see this little bump better.

(Transparency)

This, now, is a feature that appears as you increase the drive.. As you increase the drive, this thing has a longer time of flight, which means its velocity of propagation is changing with drive level. For third sound this is really a complete mystery. People saw this kind of thing early on.

In fact, when nonlinear physics first hit and people discovered solitons and the KdV equation for finding solitons, there was a big fad, a big rush, to show that your system, the one that you had been studying all your life, also had solitons in it.

People who were doing third sound took the equations for third sound -- I do not know the name for this mathematical technique -- which were quite messy, put a bunch of terms over here that would look just like the KdV equation, and they had other terms over here, and they said, "Well, let's set these to zero. Ah, we have KdV solitons in third sound."

Of course, they did that, published papers, and there was absolutely no agreement at all with the experiments. But anyhow, they got the KdV equation. That is not substitution. That would be what? "evaporation" of terms?

DR. GARRETT: Selective truncation.

(Laughter)

DR. MAYNARD: So there is no explanation for this nonlinear effect, but it is definitely a nonlinear effect, because the velocity of propagation depends on amplitude. So it is a nonlinear entity.

(Transparency)

The next thing we did was to look at how this nonlinear thing behaved in a periodic array of scatterers. Here is a third sound pulse going over a periodic array of scatterers and it looks like just the pulse coming through, with maybe some other stuff here.

But if you think about it, it cannot be just the pulse, because every scatterer reflects that pulse. I mean, it would have nowhere near this kind of amplitude after going through all those scatterers.

What happens is that it scatters back and then that gets forward scattered, and then there are all sorts of interference effects going on. But for a periodic system it all comes out coherently to give you a decent pulse. It may not come out at the same speed of sound, which is what Dr. Marston was saying.

It also has all these little things here, and these are not just simply reflections from one backscatter; it is much more complicated than that. But there is structure here.

DR. MIGLIORI: This should map onto the light transmitting through it.

DR. MAYNARD: Yes.

If you take the Fourier transform of this, you get this. This is actually sort of decent band structure. Here is a band -- here (and this should drop all the way down to a gap) -- then here is another band, a gap, and so on. We actually took other data where we got much better -- we used a pulse, but we Fourier transformed it and got nice bands and gaps for the periodic system. But that is supposed to work for the linear system.

For this experiment we increased the drive level. When you increase the drive (you normalize things when you plot them), you find that all this sort of structure goes away. When you take the Fourier transform, you do not get any band structure. What we think is happening here -- remember that for a nonlinear pulse there is this other length -- the stronger the nonlinear pulse, or the bigger the amplitude, this length gets smaller.

So even in the periodic system, when you have a strong pulse, this nonlinear length is so short that it does not really probe enough of the scatterers to even know that it is periodic, so the band structure goes away. That is consistent with the theory that was applied to the disordered system.

(Transparency)

Then we looked at the disordered nonlinear system. It is the same as we did for the periodic; it is just that the scatterers were put down in a disordered way. One of the difficulties with the

disordered systems, and with this experiment, in particular, is if you have a disordered system and you make some measurement, you might get some value, but values are going to fluctuate wildly.

In fact, even these eigenfunctions really do not have this nice exponential decay; they fluctuate above and below it.

Transmission through a disordered system fluctuates wildly. Universal conductance fluctuations are another example.

Theoreticians, when they do calculations, can get any answer you want. The only way they can get a unique answer is by doing ensemble averaging. What that means is you have some distribution function for your disorder. You make one realization, you do a calculation or you make a measurement and you get a number. Then you take another realization, get a number, and you average those. You ensemble average over a bunch of possible realizations of the disorder, but with the same sort of statistics in the disorder. That way you get a number.

When you are doing experiments, we can put down one realization of the disorder and do a measurement, and that can be kind of a big deal. But then to have to do it again and again for different realizations of the disorder gets to be a lot of work.

For the third sound experiment we had one long array of disordered scatterers and we put a number of transducers in there at different distances. We used different pairs to get different lengths, or we had different pairs at the same length but over different sections of the disorder. So we could do some ensemble averaging. That took a long time.

(Transparency)

When we got the data -- these, now, are the data for the third sound pulse over the disordered system -- at the level where the width of this pulse was about the same as the Anderson localization length we did see pretty much of a break in the plot. We think we have seen this effect that was predicted for the intermediate sized nonlinear pulse.

(Transparency)

There was another nonlinear experiment we did that involved the wire with the masses on it, but with the masses arranged periodically. This fits in with what Dr. Keolian was showing in his demonstrations.

He showed you some nice demonstrations of systems with bent tuning curves. If you have an element that is nonlinear, like a nonlinear spring that, let us say, in this case, stiffens when you go to large amplitudes, for low amplitudes the response is a function of frequency (it just gives you the usual tuning curve), but as you go to higher amplitudes, the resonant frequency shifts up and you get this bent tuning curve.

As you go up, you fall off this end; as you come down, it will jump up here, and so on.

Now, suppose you have, instead of just a mass on a spring, a string with a lot of masses on it (in fact, we had a steel wire with 24 masses on it). Now you have 24 resonant curves close together in a band.

(Transparency)

What happens when you drive this nonlinear? These things should bend over -- like this -- and you may have places where you have a lot of them bending over on top of each other. The question is, what happens in that case?

We drove this system to larger and larger amplitudes. Again, this was normal mode, this was not a pulse. This was just driving a periodic system, sweeping the frequency, and you get a band with the 24 resonant frequencies in it. Now you drive it at larger amplitudes and sweep the frequency. You can see the resonant curves are bending over and then dropping off -- bending over and dropping off -- each one independently.

As some of them start to bend over their neighbors -- it gets a little complicated -- like this, the unexpected thing is that right up here at this end of the curve it got very busy. It looks like noise, but this is not noise, because this is a plot of response as a function of frequency.

We would drive the system and it would be going cosine ωt . Then we would change the frequency a little bit and there would be a dramatic change in the amplitude. There is a little noise on here (actually, temperature fluctuations in the room). But here you sweep up in frequency and it just drops off, just as it would for a single oscillator.

But here you sweep up in frequency and there is dramatic change, almost random-looking change, in the amplitude. It does not just go up and drop off.

We zoomed in on this. We did a sweep -- I think this is a kilohertz here -- we were doing sweeps at millihertz and it just gets busier when you look in closer. It looks like some sort of fractal business going on here.

We wanted to try to understand this. These nonlinear experiments can be very difficult and challenging for an experimentalist. For example, we have this wire with the little masses on it, and the sections of the wire are doing this, when you drive it, the resonance, but there is also the other mode -- this way -- and they can be degenerate.

One thing the string can do is it can start to whirl and the little masses can start to torque like this, and it just becomes a mess. And it is at low frequency that all of this is going on. What is happening is all kinds of nonlinear things are pumping energy into low frequencies and we are getting all these other funny modes going on.

We tried to work around that with the experiment. We had a wire that had little needles on it that were poking down into little dash pots in oil, so that it would go like this and would not do so much of that -- it sort of worked.

We finally bit the bullet and did a computer simulation. To do this, it is nonlinear, so it is a long-winded computer simulation. We modeled the string as a bunch of little masses connected with springs, and then there would be a big mass that would be like the mass on the string -- big masses, little masses, and so on. Then we had nonlinear springs in there to model the finite amplitudes.

We drove that in the computer and swept the frequency and the spectra look a little different. We had a problem with the driver here. We re-did this experiment and got much better results, but we did not have quite as many data for the nonlinear drives.

It is supposed to be bigger at this end than it is at this end, this band. As you sweep the frequency, you go into the band, and you get the resonances like that. As you increase the amplitude, you can see -- right in here -- there is some busy stuff going on. There is more stuff here than just the original frequencies.

(Transparency)

We zoomed in on that in the computer simulation and we got this. Here is the resonant frequency -- here. And here is the sequence of the others. But there is all kinds of busy stuff going on here.

It reproduced in the computer simulation, which, in some ways, is cleaner. By the way, computer simulations of nonlinear things are not clean, either, because you have numerical roundoff problems that are almost as bad as noise in your experiment. So that is tough, too. But it certainly showed the same effect in the same place, so we really believed it.

We wanted to try to understand that theoretically and we wanted to take a look at what the bent tuning curves looked like; we wanted to calculate them analytically. This was a periodic system, and you can calculate periodic systems analytically. That, in fact, is what got this whole paper started.

(Transparency)

When I sat down to calculate the periodic system, it looks like this. There is a finite system of N -masses, but they are arranged periodically. It is clamped here. You know that the eigenstates are Bloch eigenstates, this nice simple form. You just want to find the driving point impedance, and from that you can calculate those transmission curves.

You actually know what the wave function looks like, you have everything here, you just calculate the impedance, it should not be hard. There it is. This is actually just the real part of the numerator. There is a denominator and there is also an imaginary part. It is that messy, but it is the closed-form analytic solution for N -masses.

One of the reasons this is so messy is this. If it did not have damping, then I would get just delta functions at the eigenfrequencies, so I had to do it with damping. When you do that, for example, you can make ω over C complex. But now you have to make the Bloch wave

number complex, too, and you have to match those up. A lot of that business is what is going on in here. It was a horrendous problem. But I got it done.

You can simplify this. I wrote it all out just to make it dramatic. But this term -- right here -- there it is, again, and there it is, again. You can simplify it by substitution.

I coded it all into the computer. You can solve this problem numerically by starting here with zero condition, multiply by 2×2 matrices, and get the impedance here, also, so you can do it numerically. I did it numerically and then I evaluated the analytic expression, and they did not agree. I went back and found that I had a sign wrong in my program.

I changed that and got agreement between the two. So that was really good.

Now I had an analytic expression for that curve.

(Transparency)

There it is. Once you have this analytic expression --

DR. MIGLIORI: Is the number of normal modes the same all the time?

DR. MAYNARD: Yes, I think that is 24 normal modes.

If I drive this at a certain frequency, there is a function that is a function of ω and a function of the speed of sound, C , and that gives me an amplitude. But the speed of the wave depends on the amplitude.

I had a program that would pick an amplitude, calculate C , and solve for the amplitude, and if they did not agree, I upped the amplitude a little bit until they agreed.

(Transparency)

When you do that, there is the result you get -- these nice bent tuning curves. The interesting thing here is that when you have this whole band of tuning curves, whole sequence of tuning curves, they can bend over in such a way that at this one frequency you can have any of 10 modes going at the same frequency.

What you might think would happen is if you sweep the frequency, you might go up one of these stable sections, it would drop off -- now, it might drop off anywhere, it may drop off way down here. Presumably, once it drops off and finds another curve, it should just move up along that curve. But it did not do that.

It would drop off to one amplitude. You would change the frequency and it would go jumping up to some other amplitude and just randomly select what mode it would go into. So this is basin-crowding, we think.

(Transparency)

This is a phase plot of the motion of the thing. When we drive it, it just goes $\cos \omega t$; there is no chaos or anything like that happening. If you plot displacement versus velocity, it is just going around in a circle here.

At some point, we say, okay, now we are going to change the frequency. Let us say we decide to change the frequency at this point and say we somehow instantly change the frequency. When you do that, what it was doing when you changed the frequency becomes the initial condition for what it is going to do after you change the frequency.

With this as initial condition, what it does after that depends on what basin of attraction it is in. It might be in a basin of attraction that takes it to this limit cycle here, or whatever you want to call it. Or it might be in a basin that takes it up to here.

If these basins of attraction are very crowded all around here -- it depends very sensitively on just when you do it. You can go to any amplitude you want.

I do not know if that is right or not, and we never dared to write this up, but if anybody has any ideas about nonlinear systems, I would be glad to talk about them.

(Transparency)

This is something new that I have not done before at the summer school: superdiffusion. What superdiffusion is, is the following. You have a disordered potential field. If you have a static potential field, then there is an eigenfunction that is localized at some position. Where it is localized depends on the actual realization of the disorder. So you may drive it at some eigenfrequency and you will get some localization at a particular site that depends on that realization of the disorder.

But now suppose the potential field is time-dependent, so the disorder changes, but suppose the disorder changes in such a way that you take just this realization of disorder and shift it over. (I emphasize that, but I have got some big spacings here -- there are the big spacings when you shift it over.)

If you do that slowly, the Anderson localization shifts over with it.

Anderson's paper was "The Absence of Diffusion in Disorderd Systems." If you have disorder, the electron gets localized, so it cannot diffuse through the system. But if you have a time-dependent potential field and a really pathological one, you could have the localized state zipping along and you can have that electron essentially moving faster than diffusion. That is what is called superdiffusion.

You do not do it by having the realization just nicely moving over; that does not happen. Apparently, according to the theoreticians, if you have a potential field that is time-dependent, not completely random but it has some correlation in it, they claim you will get a transfer of the electron wave packet at a faster rate than diffusion.

If it were diffusing, it would go as the square root of time. In superdiffusion it goes as time to some three-halves power, or something like that.

We wanted to see that in an experiment. One way you might think of doing these experiments -- you read Anderson's paper and you see this potential field that looks like that

-- one of the first things that will pop into your head is that, well, I can have a channel of water and I can have the bottom look like this, so there is the speed of sound changing as a function of position (because the speed of sound depends on the depth). This is the kind of thing that might pop into your head for doing Anderson localization, looking at surface waves in a tank or channel, whose bottom varies like that, and people actually tried to do that.

It did not work, because water is too lossy. You need long-range phase coherence for it to work.

FACULTY: Use mercury.

DR. MAYNARD: Yes, well, you could.

In this experiment you might think about doing the localization experiment and have the time-dependent potential by having, say, some sort of rubber membrane on the bottom and you drive it up and down with motors. So there you have your time-varying potential.

But the problem with that is that if you drive the bottom up and down, that motion is going to generate a wave on the surface, and that is not allowed. Your change in potential cannot generate wave function.

In a solid, time-dependent ion motion cannot generate electrons. In fact, in the theory this is also a problem. In the theory, you have to maintain the normalization of the wave function. There were some back-and-forth letters in *Physics Review Letters* about that.

In an experiment that is a problem. You have to be able to change the potential field without generating a wave.

(Transparency)

Having studied with Izzy Rudnick, you know what to do. For a one-dimensional wave medium, we use a strip of metal. It is just hanging loose. When you shake it, you get bending waves in that strip, so that is a one-dimensional wave medium. Along this strip we put little piezoelectric chips in the form of a bimorph.

When you put a potential on it, the chip on top gets longer, the one on the bottom gets shorter, so it bends the thing that way. When you do that, the thing gets a lot stiffer -- I mean, this is really exaggerated -- but like that, it flops down. I can make it stiffer by bending it in the other direction. It is like a venetian blind.

By bending it like this, we can change the speed of sound in this strip. Our bending waves are going down this way, but we are changing the speed by bending it the other way. To low order, doing this does not generate any waves in that direction. In fact, we tested that and it works fine.

Now we have a beryllium copper strip with these piezoelectric chips in it and there are, I think, 40 of these bimorphs along the strip (it is about 8 feet long). Then there is a shaker to put the transverse waves in the strip. There is a computer bus here, with cards in it and some high-

voltage drivers, so we can use the computer to -- in real time we can change the voltage on the bimorphs.

The speed of the wave in the strip is quite slow, maybe a kilohertz is the top frequency you would have. You could certainly change the speed of sound as a function of position much faster than that with the computer setup.

So now we can have a time-dependent potential field and look at the behavior of waves in the system. We are just now finishing the construction of this. There are 40 circuits here that had to be built -- a lot of work. We are just now getting this on line and, hopefully, we can make the first experimental observation of superdiffusion.

These experiments involved literally balls and springs and serious contemporary physics and they are a lot of fun. Thank you.

DR. GARRETT: On that last experiment, do you have to put the dc bias on all 40 of those so that there is some initial minor amount of curvature, because the propagation properties of the flat, uncurved strip are so different from the strip that has any curvature in it whatsoever?

DR. MAYNARD: It is not that different. We worked on this and we did test sections and everything. You get a nice continuous change in the spectrum; we do not bend it very much. In fact, with those bimorphs you cannot get all that much -- we looked in your engineering book, as a matter of fact, for what happens when you bend a plate like that. It seemed to work fine.

MR. BENNETT: The Anderson length, what is the definition of that?

DR. MAYNARD: That is a good question. In fact, that is something we worried about for a long time. The eigenfunctions really do not just have this nice exponential envelope. That happens only with probability one. A real eigenfunction can fluctuate up and down.

One way to understand what is going on is to look in Marshall Luban's paper. He explains how, in your head, you can deal with those ups and downs.

In this paper that I wrote there is another quantitative way of looking at that. You use a thing called the "participation ratio." You make a plot of this as a function of distance down the system. It will fluctuate, but if it is a periodic system, if there is no localization, as you plot this thing, it will just keep going up. It will fluctuate, but it will just keep going up.

If there is localization, this thing will go up, but then it will level off. Even if there are fluctuations, you can see that happening. That is really the proper way of looking at localization. The place where it levels off, that distance is the Anderson localization. That is a quantitative way of doing it.

DR. MARSTON: A demonstration that you may want to think about in regard to this experiment you are developing is to take your tape measure and flex it and look at what happens to the orthogonal curvature in the region of the bend. Just try it and see what happens.

That actually was a paradox. The response in that case was proposed to Love (in Love's book on elasticity) and he did not solve it. I think some 50 years later somebody wrote a simple discussion of how to understand the response.

DR. MAYNARD: We are not bending it very much.

DR. SWIFT: Back a few viewgraphs when you had the family of 24 tuning curves that were bending over, it is this experiment that I want to ask about. It seems when you have something that is that nasty, that as you change frequency in your experiment, you really have to ask yourself what does the frequency synthesizer do when it changes frequency? Does it wait for a zero crossing and then begin with the new frequency, or does it wait for a peak and then begin, things like that?

DR. MAYNARD: The actual structure here will not reproduce just because of that. But the existence of that structure is really robust.

DR. SWIFT: So if you use a frequency synthesizer and you step through with a different engine, the 12th digit instead of the 11th digit and get different --

DR. MAYNARD: But you could also get this coming down in frequency as well.

DR. SWIFT: Everything is a little different if you are coming down up there in the middle.

DR. MAYNARD: But this little business going on at this end -- oh, there is another thing, too. It always happens at this end, but you will notice that the big peaks are at that end, as predicted by the theory.

DR. MIGLIORI: So how does the specific heat go?

DR. MAYNARD: For what?

DR. MIGLIORI: The number of degrees of freedom of the system of something like this.

DR. MAYNARD: I do not know. I have not thought about that at all.

DR. MIGLIORI: It seems hard to believe you get more degrees of freedom for real. At least at some level you ought to be able to configure a nonlinear system in which the degrees of freedom are really fixed.

DR. MAYNARD: Wait, I am not sure what you are worried about. These are not really degrees of freedom here. This is just the amplitude of the system as a function of frequency.

DR. MIGLIORI: Yes, but that is the sort of thing that generates specific heat in a solid, because you are really looking at the $D \log \omega$ -- something -- like pressure.

DR. MAYNARD: I think that ω is an ω for a normal mode. In this nonlinear system that is no longer a valid concept, I think. That is why I am having a little trouble.

DR. MIGLIORI: Okay, but heat is a measurable quantity in the system, and somehow, somewhere, I am kind of searching for a conservation law that says that specific heat at high temperatures is dependent on a number of degrees of freedom.

DR. MAYNARD: I have a feeling that in that kind of a system the nonlinear coupling of the thing is so much weaker than this that it is a different thing altogether. You have some nonlinear coupling between modes, you know, so you thermalize, but you can still identify the modes.

In a nonlinear system I do not know what it means to try to identify modes. It is almost like what Dr. Keolian was saying the other day.

SONOLUMINESCENCE

ANTHONY ATCHLEY
NAVAL POSTGRADUATE SCHOOL

DR. ATCHLEY: I want to give some encouragement to everyone out there. If you take Dr. Migliori's talk, just to pick one, there are very few people in this room -- or in the world -- who could give that talk. If you look at the combination of experience and expertise, the broad range of knowledge, the being at the right place at the right time, the chances of people out there giving that talk are almost zero, almost as much error as he has in his results.

I, on the other hand, stand before you giving a talk that any of you could give. The field of sonoluminescence at this point in time is so small and so young that in a day you could go to a library and pick out 75 percent of the important references. You could read them in a day and give this lecture at night. It is rare to find such a field these days. Even though I'm giving this talk, whatever I say about my own work in this field is really of no consequence for the purposes of this talk (and people tell me that all the time).

(Laughter)

In the way of an outline, most of this talk is a review of experimental results, primarily because there really is no theory that I can teach you about sonoluminescence. We just do not know enough about it. There are theories being developed but, as far as I know, there has been no theory that has predicted anything that had not previously been measured. There are no testable predictions from any model, so I cannot teach you what the right answer is.

What I can tell you a little bit about is the background of sonoluminescence; that is, the mechanics of getting a bubble to emit light. So I will spend a little bit of time talking about acoustic levitation. We will spend a little bit of time talking about bubble dynamics, just to give you a feel for why you might expect, under the right conditions, that something interesting could happen.

Then, I intend to talk about experimental observations, just to give you an indication of some of the startling results in this field. I will also talk about what people think is going on, but the main conclusion is that nobody really knows what is going on yet.

Let's talk about bubble dynamics. Suppose I have a flask containing water and I inject a bubble into the flask using a hypodermic syringe. What do you think is going to happen to the bubble? It is probably going to rise up. It could dissolve; but if the water is saturated the bubble is probably going to rise up from buoyancy. Therefore, if you want to try to study something having to do with that bubble you have a problem. You put the bubble in, but it is not going to stay there. Why does it rise? What force?

MR. WORLIKAR: Buoyancy.

DR. ATCHLEY: Buoyancy, right. There is a force on that bubble in this liquid that causes it to rise. Therefore, to study this bubble, you need some way of counteracting that buoyancy force. One way to do that is through a process called acoustic levitation.

(Transparency)

I have put a simplified explanation of acoustic levitation on one page. This is not the best kind of transparency for me to present, because it has lots of equations, but here goes. Let us assume there are only two forces acting on a bubble in a liquid in the presence of a sound field. One is buoyancy and the other is an acoustic force. Suppose we want to be able to suspend this bubble at a fixed point in that liquid and have it not move. To accomplish this, the buoyancy force has to be counteracted by an acoustic force, right? [Refer to Eq. (1).] What is the acoustic force on a bubble? How do you derive it? What fundamental equation of acoustics?

MR. WORLIKAR: Euler's equation.

DR. ATCHLEY: Euler's equation, or you linearize it and make it a simple force equation, Eq. (2). Equation (2) is Newton's second law per volume. The force per unit volume exerted upon an object in a sound field is equal to the negative of the pressure gradient. Because the pressure gradient changes with time, this force fluctuates with time. If this bubble is to stay fixed for all time, there must be some time-averaged non-zero force. So let us look at the time average Eq. (2).

The pressure in the liquid, Eq. (4), is the sum of the mean pressure plus some pressure due to a standing wave of amplitude P_A . The acoustic pressure varies spatially, say, as cosine of kz and has a sinusoidal time dependence. That is one part of the answer.

The volume of the bubble, Eq. (5) is composed of an ambient or mean volume -- that is, in the absence of sound and a fluctuating part. The fluctuating part is caused by the presence of sound, when the pressure increases or decreases, that volume changes. It should change in some proportion to the driving pressure and it is also going to depend on the frequency. This term is the acoustically induced change in volume. In Eq. (5), it is expressed in a form appropriate for a bubble below driven resonance: When the pressure goes up, the bubble shrinks and when the pressure goes down, the bubble expands. This is the origin of the minus sign in Eq. (5).

Why is the relationship between the acoustic volume change and acoustic pressure frequency dependent? Think of an harmonic oscillator. Below resonance, the displacement of the driven end of a spring and the displacement of the mass are in phase. Above resonance, the two are out of phase. This explains the frequency dependence.

Substituting Eq. (5) and the gradient of Eq. (4) into Eq. (2) yields Eq. (6). Equation (7) is just the time average of Eq. (6). It represents the force that traps -- or levitates -- the bubble. It is a dc time-averaged acoustic force.

How does it behave? The graph shows two curves. It is oriented the way it is to mimic what would happen in a resonator that is mounted vertically, with the sound field axis up and down. The vertical axis represents some arbitrary position in a standing wave. The solid line represents the acoustic pressure amplitude as a function of position at some particular time. The dashed line represents the time-averaged acoustic force as a function of position. What does it tell us? Suppose I put a bubble at $kx = 50$. What is the force at that point? Positive or negative? It is negative, right? Which way does the bubble want to move? Down, right? Negative force means it is going to move down. The acoustic force is going to push it toward a pressure anti-node. Suppose I put it at $kx = 200$. Which way will it move? The force is positive; it is going to move up -- again toward a pressure anti-node. As it moves it gains momentum and overshoots the anti-node. What happens to it? It comes back, because the sign of the force has changed, so it comes back to the anti-node. That is the point of stability for this bubble. This is acoustic levitation simplified.

Notice that the sign of the acoustic force is tied to the relationship between the volume change and acoustic pressure. If the bubble is driven above resonance, then when the pressure goes up, the bubble does not shrink. It expands. Its motion is out of phase -- which means, this dashed curve flips over. Instead of a bubble going to a pressure anti-node, it is attracted to the pressure node. So where the levitation point is, or which way it goes, depends on this functional dependence between the volume change of the bubble and the acoustic pressure and the frequency, at the simplest level.

Now that we understand the acoustic force, let us return to buoyancy. Buoyancy is given by the product of the mass of the liquid displaced by the bubble times g , Eq. (3). This force also depends on the volume of the bubble, Eq. (5). If you time-average Eq. (5) it turns out the acoustic part disappears and the time-averaged buoyancy force is the mass of the water displaced by the bubble times g , Eq. (9). If you equate the time-averaged acoustic force to the time-averaged buoyancy force, you arrive at Eqs. (10) and (11), which are equivalent. These equations allow you to determine the equilibrium position of a levitated bubble.

The main point to take away from this discussion is that you can counteract the buoyancy force acting on a bubble by applying an acoustic field. Once you do that, you can put a bubble in this liquid, and it will not rise; it will stay put. Not only will it stay put but you know where it is going to be and can study it for a long time.

MR. WORLIKAR: Does it stay at one point or it oscillates at some point?

DR. ATCHLEY: There is some time-dependent motion; it does oscillate about the equilibrium point. But usually, for the bubbles I want to talk about, that effect seems to be unimportant.

(Transparency)

All right, so at least now we know how to keep a bubble in place. Now let us worry about the response of the bubble in more detail. If you study the field of bubble dynamics, there is an equation you run across almost from day one called the Rayleigh-Plesset equation that discusses bubble dynamics.

(Transparency)

I want to go through a simplified derivation of that equation just so you have a feel for where it comes from. Assume that a bubble of radius R is immersed in an infinite, incompressible fluid. We will allow the size of this bubble to change and apply conservation of energy. The time rate of change of kinetic energy has to equal the negative of the rate of change of potential energy. All we have to do is find kinetic and potential energies for this bubble.

Kinetic energy is just equal to the mass of the liquid that is moving -- density times volume -- times velocity squared. But because the velocity of the fluid depends on distance from the bubble, one has to integrate the mass times the velocity throughout the volume as in Eq. (13). If the fluid is incompressible, the velocity of the fluid u at some distance r from the origin can be related to the velocity of the fluid at the bubble wall through application of mass conservation, Eq. (14) and (15). (Capital R denotes bubble radius, lower case r denotes distance from the origin.) Upon substitution of Eq. (15) for $u(r)$, Eq. (13) can be integrated.

(Transparency)

The details of this integration are shown in Eqs. (16) - (19). It is typical in the literature to use \dot{R} instead of using $u(R)$ for the bubble wall velocity. Eq. (19) is written in the way it is to make a point. Recall that Professor Garrett talked about bubble oscillations and he said the effective mass of a bubble -- at low frequencies, anyway -- is equal to three times the mass of the fluid displaced by the bubble. The fluid displaced by the bubble is $\frac{4}{3}\pi R^3$ times the density of water. So I have written the bracketed term like that; that term is the mass of water displaced by the bubble and, to emphasize the water density, I have written it as ρ_L . Three times the bracketed term is three times the mass of the water displaced; it is just the effective mass of the bubble. Therefore, kinetic energy can be written, of course, in the form $\frac{1}{2}m\dot{R}^2$, where \dot{R} is the velocity of the bubble wall, and m is effective mass. We could have jumped right there. I could have just said kinetic energy is $\frac{1}{2}m\dot{R}^2$. Professor Garrett told us what m was. Problem solved.

Now, we have to take time derivatives, so we go back to this form and use most all of Professor Garrett's important rules of mathematics, except Taylor expansion, and you get Eq. (23) for the time derivative of the kinetic energy. So much for kinetic energy.

(Transparency)

Potential energy: Suppose we have a bubble in this fluid, and even though I have drawn the fluid to fit on here, the extent of this fluid is infinite. Let the pressure outside the fluid be p_∞ .

infinity, and let the pressure at the bubble wall, but just in the fluid, be p_{∞} . If these two pressures differ, the bubble will change volume. The rate of change in potential energy is just the pressure difference across that fluid times the rate of change of the volume of the bubble, Eq. (24). You take the time derivative dV/dt and you get Eq. (25). That is what the rate of change of potential energy is. And we are done!

(Transparency)

Setting the time derivative of the kinetic energy equal to minus the time derivative of the potential energy, Eq. (26), results in Eq. (27), which is called the Rayleigh-Plesset equation.

This equation is the starting point for any discussion of bubble dynamics. In 1917, Rayleigh treated a case when the right hand side of the Rayleigh-Plesset equation is constant. He treated the following problem: Suppose a spherical void of radius R is instantaneously created in an infinite, incompressible fluid. What happens to the radius of the void as a function of time? So you have the case where, boom, you make a bubble with nothing inside, the initial velocity of the wall is zero. There is potential energy stored in the fluid. Now you let the bubble wall go. What happens? The answer it does that it collapses very rapidly. Rayleigh found extremely high collapse rates.

The reason Rayleigh worked on this problem was to find out why fast ships slowed down so dramatically after a few days at sea. There was a ship that left England -- my brain is frozen, I cannot remember the name of it -- trying to beat the trans-Atlantic speed record. It started off the first day and it was just going like gangbusters. By the second, third, and fourth day it was practically stopped in the water. They hauled it out and the propeller was basically gone.

The reason was that the propeller was going so quickly that the pressure in the liquid was below atmospheric pressure, below ambient pressure. This reduced pressure created vapor cavities near the propeller. When the cavities collapsed, they did so on the surface of the propeller with such velocity and force it eroded the propeller. Each little bubble was like a jack hammer, an arc welder, that pitted the propeller. After a few days the propeller was practically useless. So that is what Rayleigh did. Plesset got hooked up to it when he started applying acoustics to the problem.

If one wants to do bubble dynamics calculations that are at all realistic, some refinements are required. One refinement is defining the pressure at infinity. It is equal to the ambient pressure in the fluid plus the acoustic pressure far away from the bubble wall. The pressure in the bubble wall on the liquid side is equal to the pressure inside the bubble, which is a combination of the noncondensable gas and the vapor. There is a correction for surface tension and viscous stress across the bubble wall. There was a nice paper written in the 1950s by Poritsky addressing the apparent paradox of how there are viscous interactions in the bubble when all the viscous terms apparently drop out of the relevant equations. I am not going to try to derive it, but there is a term

you have to put in to account for viscous stresses across the bubble wall as you go from inside the bubble to the outside.

Professor Marston can probably answer that better.

DR. MARSTON: The basic problem is you have to think about the rate of strain tensor in spherical coordinates for spherical flow and it is not as obvious.

DR. ATCHLEY: Which is why I did not do it. I can give you the reference.

DR. MIGLIORI: Is this only a surface effect?

DR. MARSTON: No. The dissipation is in the volume. A good place to see those kinds of expressions written out is in the appendix of Batchelor's book on fluid mechanics.

(Transparency)

DR. ATCHLEY: You also need to include compressibility. Compressibility does two things -- this is compressibility of the fluid. First, when the bubble expands, the fluid does not move instantaneously, so the expression we use for the velocity of the fluid as a function of radius, Eq. (15), is not exactly right, because there are time delays. Also, when the bubble is at its biggest and starts to collapse, then some of the stored energy out there in the fluid does not necessarily come back to the bubble. Some of it is radiated away in the form of sound. These are the two biggest compressibility effects.

To put compressibility into the derivation would require much more detail than I want to get into, and it would really put you all to sleep. I'll just say that there is another form of the Rayleigh-Plesset equation, called the Keller equation, that takes compressibility into account. There are corrections that have the Mach number of the bubble wall built into them, and there is also a radiation term.

The Keller equation is the one that you should really solve, if you really want to start investigating bubble dynamics.

MR. WANG: What is the wavelength compared to the bubble size?

DR. ATCHLEY: In this derivation, the wavelength is large compared to the bubble. kr is a small number.

(Transparency)

All right, let us talk about the Keller equation. It looks a little bit nasty, but it is easy enough to solve. It is an ordinary differential equation, and I can solve that using MATLAB or Mathematica. I have solved the Keller equation, Eq. (30), for three different values of acoustic drive amplitude. The bubble radius is 5 microns. The acoustic pressure amplitudes are 0.05, 0.5, and 1.25 atm. The ambient pressure in the liquid is atmospheric and the bubble is driven at 25 kHz. This upper graph on Transparency #6, Fig. 6-a, shows the acoustic pressure as a function of time. The unit of time is microseconds. The range is 60 to 160. The reason that the time axis

does not start at zero is to allow time for transients to die out. Two or three oscillation cycles of the bubble are shown.

At low-drive amplitudes, what would you expect to happen? You have this bubble, you are driving it gently, what would you expect to happen to the bubble radius as a function of time? How should it behave? As the pressure goes up, the bubble ought to shrink. If the pressure goes down, the bubble ought to expand. If the drive is sinusoidal, the response should be nearly sinusoidal. This is the behavior shown in Fig. 6-b, which is the 0.05 atm drive case. It behaves as it ought to; the response is linear.

Now, let us crank the drive up by a factor of 10. What happens? You can see clear distortion in response to the bubble. At the low-drive level (Fig. 6-b), the maximum and minimum deviations in radius from the equilibrium value are the same in magnitude. This is what one would expect from a linear response. But when the bubble is driven harder (Fig. 6-c), the largest radius deviates more from the equilibrium value than does the minimum radius. The radius-time curve is distorted. The bubble does not oscillate sinusoidally.

Now, crank the drive up even more, to 1.25 atm. The response is shown in Fig. 6-d. Now, the bubble behavior is really odd. Instead of oscillating modestly about the equilibrium radius, it gets big and then, boom, it collapses very quickly to a very small size. In Fig. 6-b, the fluctuation in radius is roughly 0.05 microns out of 5. In Fig. 6-c, the fluctuation is roughly half a micron, a 10-percent change in radius. But in Fig. 6-d, the bubble starts at 5 microns; it goes up to approximately 35 microns and then slams down to a very small size.

There is one more aspect of the bubble behavior to be pointed out. The bubble spends a large fraction of each cycle undergoing small amplitude oscillations around its equilibrium radius. Also, the period of these oscillations is much shorter than the period of the drive. The bubble does not exhibit the typical sinusoidal behavior where it spends most of its time away from equilibrium size. In Fig. 6-d the bubble spends a lot of time near equilibrium.

Does anybody know the cause of these small amplitude oscillations? The periodicity of those little bumps matches the resonance frequency of the bubble. When it collapses, the bubble sits there and free decays for a little while; it is not really being driven. It is being imploded due to the inertia of the fluid. When it gets really big, it stores a lot of energy in the fluid and then, when the bubble begins to collapse, the inertia of that fluid just presses it in, and it crunches to a small size. The bubble then oscillates at its resonance frequency, trying to regain its wits. About the time it settles out, oops, the drive goes again and the whole cycle repeats.

DR. CHAMBERS: When it gets down to a small size, does it ever become a liquid? I know it does not go to zero radius, but does the gas become a liquid?

DR. ATCHLEY: That is a very good question. What happens in the interior of the bubble is a big unknown.

DR. MAYNARD: What would happen if you drove the bubble at its resonance?

DR. ATCHLEY: It would go crazy. It would not stay -- first off, levitating a bubble at resonance is a question.

DR. MAYNARD: Well, of course, yes.

DR. ATCHLEY: Also, it would probably grow so much from diffusion it would not be stable and you would not get this behavior. These bubbles are all small compared to resonance size.

DR. GARRETT: You should point out that that effect is known as "Disney dynamics," because whenever you see a cartoon where somebody gets whacked, he always goes "boo-boo-boo-boo."

MR. LANDSBERGER: The pressure that is driving the bubble into oscillation, is that the same pressure that is levitating the bubble?

DR. ATCHLEY: In practice, yes.

STUDENT: What is the mean pressure?

DR. ATCHLEY: One atmosphere.

MR. WORLIKAR: You are driving at an amplitude high enough for the pressure to be negative?

DR. ATCHLEY: Yes, when the pressure is peak negative, the instantaneous pressure is minus 0.2 atm. The liquid is under tension, just like you if would take a piece of steel and stretch it.

If you prepare a liquid very carefully, then you can approach theoretical tensile strengths in the liquid, and for water it turns out to be a couple thousands of atmospheres just because of the attraction of the molecules. When the pressure go below vapor pressure, bubbles do not automatically form. It is exactly analogous to superheating or supersaturating in chemistry. You can supersaturate a solution of sugar water and as long as you keep any nucleation sites out of that sample it can contain more sugar in solution than it should. But as soon as you introduce a site for that sugar in solution to precipitate out, then it does so rapidly.

The same thing happens in cavitation. In order for a bubble to form under ordinary conditions, there has to be a nucleation site, and the most common nucleation sites are dirt. If you filter the dirt out, then you can easily attain tensile strengths of 5 atm. If you do a better job of filtering, you can get to 15 atm. Mo Greenspan was able to get to 200 atm, and it turns out the nucleation site for those experiments was not dirt, but bubbles generated by cosmic rays and neutrons that had been through the liquid. He did this beautiful experiment where he took some clean water and put it close to a neutron source. He showed that the neutrons going through the liquid would cause microscopic bubbles to be formed and that would rupture the liquid. It was a bubble chamber.

Back to bubble dynamics. Keep in mind all I have told you is what happens to the bubble wall. I have not said anything about what happens inside the bubble. Going from the liquid into the bubble is tricky. The simplest thing you can do is forget about there being any water vapor or liquid vapor in there and just assume an ideal gas that obeys the adiabatic gas law inside the bubble. So PV^{γ} is a constant.

(Transparency)

Let us see what happens inside the bubble now for the same drive conditions as Fig. 6-d, a 5-micron bubble driven at 25 kHz at a pressure amplitude of 1.25 atm. The R--T curve is reproduced in Fig. 7-a. Fig. 7-b shows the derivative of the R--T curve, or the velocity of the bubble wall versus time. Keep in mind that what I am about to say is probably garbage, because when you get large velocities this Keller equation is not necessarily valid. However, it is indicative of what might happen in reality.

Notice that when you get the big collapse and the rebound, that you get large negative velocities and large positive velocities. And I will just say large; do not take these numbers with any faith. But that is a pretty impressive velocity, 1000 meters/second. Something interesting could happen under those conditions.

What about the pressure inside the bubble? (See Fig. 7-c.) If we assume that gas in the interior of the bubble obeys the adiabatic ideal gas law, then at the time of the big collapses, there are big pressure spikes. I will just point out the exponent on the pressure axis is 9. Again, the actual value is probably garbage, but that is an impressive number. The point is, once again, something interesting could happen. Even if the numbers are off by a couple of orders of magnitude, something interesting could still happen.

The temperature inside the bubble (Fig. 7-d.) is also spiky, and also indicates that something interesting could happen inside the bubble. Keep in mind this is a very simple equation, very simple assumptions, reality is certainly different, but if you were looking for something interesting to happen with a bubble you should look at it during the collapses.

(Transparency)

This is a graph is some of Felipe Gaitan's original work, published in JASA. It shows the acoustic pressure drive (amplitude is a function of time) in the upper curve. The middle curve shows a measure of the bubble size determined from a light-scattering technique. It similar to the R--T curves shown previously. The lower trace is the output of a photomultiplier tube. Every time the bubble collapses a burst of light is produced cycle after cycle. These results were the first observations of picosecond sonoluminescence from a single levitated bubble.

(Transparency)

For historical purposes, as far as I know, the first paper written on single bubble sonoluminescence was in the proceedings of a conference that David Blackstock and Mark

Hamilton put together in Austin, Texas in 1990, the 12th International Symposium on Nonlinear Acoustics.

So far, we have examined a bubble being levitated and driven acoustically. We have examined how the radius behaves as a function of time. We have looked inside the bubble and found out that simple equations predict huge pressures and temperatures, too big to be believed, almost certainly, but indicative that something interesting could happen. We have also seen that if you observe a bubble being driven under these conditions, something, in fact, does happen. The bubble emits light. After the break, we turn attention to some of the properties of this light.

DR. ATCHLEY: Why was Felipe's discovery of single bubble sonoluminescence so important? After all, people had observed sonoluminescence 60 or more years ago, when they started building high-intensity transducers for sonar. Pretty soon they determined that bubbles were the source of the light. From, say, the 1930s through the 1980s, people had used sonoluminescence as an indicator for a number of things. For instance, when I did my dissertation on cavitation thresholds, we used sonoluminescence as an indication the liquid had ruptured and there had been a violent collapse of the bubble.

The literature on sonoluminescence is rich and goes back a long time. But nobody, with a few exceptions, studied the properties of the light itself. There had been measures of the spectrum and attempts to measure the duration of the emission. However, the basic trouble was the process was too random in both space and time to really figure out how to measure the properties in any reasonable way. So up until the time that Felipe did his work with the single levitated bubble, so that one could know where and when the emission would occur, there had not been a lot of investigation of the properties of sonoluminescence. But that all changed.

(Transparency)

The method that Felipe used to generate sonoluminescence is the following. A low amplitude, standing wave is setup in a glass walled, rectangular cross section, container that was filled with filtered, degassed water. A volume of air is injected into the water with a hypodermic syringe. The majority of this gas rises out of the water. However, a small portion is trapped at the pressure anti-node. This small bubble is perhaps 50 microns in radius, easily seen with the unaided eye.

These conditions are similar to the low drive conditions discussed in connection with the Keller equation earlier. The bubble is levitated near the pressure antinode and oscillates in a well-behaved manner. This condition is indicated in the lower region of the illustration shown in the transparency on page 209 of Vol. II.

As the bubble is driven with higher and higher amplitudes, its radial oscillations become larger and larger. At some point an instability sets in. The bubble is dancing all around. As the drive amplitude is increased further, the dancing becomes much more violent and the bubble

appears to shed smaller bubbles, shedding mass. Upon further increase in drive amplitude, the dancing motion ceases completely. What one has now is a very small -- perhaps a few micron radius -- very stable bubble. At this point the bubble glows. As the amplitude is increased more and more, the bubble glows brighter and brighter. There is, however, an upper limit to the drive amplitude which can be used. If the bubble is driven above this level, it disappears.

Well, here you have a bubble driven, maybe, at 9/10 atm. Here you have got a bubble driven 20 percent higher. Here you have got no photons. Here you have got about a million. With that small change in drive amplitude, some interesting physics happens in that 20-percent change in drive. That got Seth Putterman's group interested.

A lot of what I am going to show for the next few minutes are results from Seth's group, because they really grabbed hold of this and went with it.

MS. GRANT: How did they know what frequencies to drive it at?

DR. ATCHLEY: You want to set up a standing wave in your resonator, so the frequency depends on the size of the resonator.

MS. GRANT: And not on the bubble?

DR. ATCHLEY: Not on the bubble. The bubble, in terms of the drive frequency, is not important. If you have a spherical flask, you want to set up a radial mode. If you have a rectangular flask, you want to set up a mode that has pressure variations along the axis and also in the other two directions to trap the bubble.

I did not mention it, but if you had a plane-wave mode -- I guess you cannot do that easily with a glass wall -- but if you had a plane-wave mode, the bubble might be constrained in one direction but there is no force to trap it transversely, so you have got to have variations in pressure in all three directions to have a stable point.

(Transparency)

One thing Brad Barber did was to look at the distribution of pulse heights for many, many cycles. They all seemed to have about the same amplitude. He also found that if you drive the bubble harder -- this is the acoustic pressure amplitude, this is the number of photons per burst, and the exponent is 5 -- as you drive the bubble harder, you get more and more light.

They made some other measurements, which I will skip over. The really interesting stuff came about when someone asked, "What is the duration of the flash?"

(Transparency)

The UCLA group got a fast photomultiplier tube that had a 2 nsec rise time and they shined sonoluminescence on it and got the output as a function of time. That is shown in this panel.

Then they took a 34 pico-second laser and hit the same PMT and they got this graph. If you look at them, they are the same, which means that you are measuring the response of the tube and not the duration of the flash. So what do you do? You get a faster tube.

They got a tube that was 10 times faster and repeated the same measurement. There is sonoluminescence. There is a 34 ps laser pulser. Again, the curves are virtually indistinguishable, which tells you that you are still measuring the response time of the system and not the duration of the flash.

(Transparency)

Brad called Hamamatsu and got the fastest photomultiplier tube they make, 170 ps rise time, and repeated the same measurements. This curve is the response from sonoluminescence and this curve is the response from the 34 ps laser pulse. Forget the differences in the turn off: That is not interesting.

It turns out the sonoluminescence turns off faster, because when the laser pulser turns off, it rises sharply, it turns off, and there is a little bump on the way down, and that is the cause of the difference in the turn off.

The interesting part of the measurement is in the rise time. They analyzed the rise time and what they found was the following.

Through all sorts of fancy processing they said that the rise time of the pulse when the PMT is exposed to sonoluminescence was 176 ps. The rise time when they exposed it to the laser pulser was 172 ps. From this 4 ps difference in the rise times they inferred that the sonoluminescence pulsation was 50 ps or smaller. But I am telling you now, that is a big extrapolation. A 4 ps difference for a rise time measurement is not much, and extrapolating the duration from that is uncertain, but it was the best they could do.

But you say there are techniques out there for measuring very, very short pulses. Why don't you use one of those? The answer is, yes, you can get really fast photodiodes, but they are used for short-pulse lasers that have a lot more photons than the bubble emits and, also, the spectral sensitivity is usually different. So even though there are fast detectors out there, they do not work for light from this kind of bubble in this spectral range. This is one of the frustrations with sonoluminescence; it is bright enough to see but it is too dim to get really good measurements on.

The report of short duration got people really going. Here is a classical acoustic system, table-top science, 20 kHz drive, water, beaker, H.P. function generator generating light with a duration so short you could not use state-of-the-art optical detectors to measure it.

(Transparency)

The UCLA group also found that if you look at the synchronicity of the pulses, that is, the timing between one and another and another, that it could be extremely stable. They found that jitter from pulse to pulse was also on the order of 50 ps.

To put this in perspective, suppose we took our sonoluminescence apparatus and put it on the top of this building and we could look down from space on Pacific Grove. Here we are in

P.G. This picture shows 10 pulses of sonoluminescence emanating spherically from this container.

At a 20 kHz drive, to go from here to Merced, which is closer to Yosemite or Fresno, takes about half a millisecond. If you could measure the spacing between these bursts, it would be 15 km. The thickness of the line would so small I cannot possibly draw it on the same scale. A 50 ps burst of light is about a centimeter and a half in space, finger size.

So you have these 15-km-spaced lines drawn 1.5 cm wide and the precision from line to line is also about 1.5 cm. That is a pretty stable and remarkably simply acoustic system.

I ganged up with some guys at Livermore and we tried to use fast streak cameras to measure the duration of the pulse, because the result from the PMT was not a direct measurement of the duration. So we took one of these fast streak cameras and tried the same thing.

(Transparency)

Here is what we get when we look at a single flash from a sonoluminescent bubble. Here is what we get when we look at a 17-psec laser pulse. We saw this and we were just on cloud nine until we realized the way we made the measurement was to have this flask here --

(Transparency)

-- I do not know if you can see this, but this is a hypodermic needle and this is an optical fiber that goes into the liquid close to the bubble, so we piped the light to the streak camera through an optical fiber to avoid mirrors and all sorts of stuff.

(Transparency)

It turns out the dispersion in the fiber should have been longer than that, so our elation deflated rapidly. I show this just to point out that measurement of duration is extremely difficult and there has still not been, as far as I know, any direct measurement of the duration of the pulse.

DR. KEOLIAN: Is there any chance, if you had an optical soliton, that would sharpen it up in the fiber?

DR. ATCHLEY: There are different dispersions in fibers. You can do pulse stretching or pulse compression. We did not think about solitons. We thought about dispersion in the fiber and we could not figure it out.

DR. KEOLIAN: The pulse is so short that you may not have had a big enough amplitude.

DR. ATCHLEY: It could be, but we sort of gave up. It took us six months to get this far. Getting that was hard enough. The stuff they did at UCLA was to average a bunch of the pulses. We tried to do it on a single shot. There are all sorts of problems with these cameras. It is a terrible system.

Brad Barber spent time in Japan at Hamamatsu with their fastest cameras, image intensifiers, streak cameras, and he could not get anything, either. It is tough.

DR. MIGLIORI: How many photons come off --

DR. ATCHLEY: Roughly a million.

DR. MIGLIORI: So you got 30 mW as the peak power level?

DR. ATCHLEY: I guess it comes out to be something like that, yes.

DR. CHAMBERS: You mentioned the stability, how quickly they get impeded and how precisely they were impeded. What is the stability of the driving PZT transducer? I have heard that the system is more stable --

DR. ATCHLEY: It is more stable than an H.P. oscillator -- but that is a tricky question, because the Q of the system is pretty high, so the H.P. oscillator could be fluctuating over some time, but when you have got a Q of 1000 or 10,000, I do not know.

All right, so the duration seems to be short. What about the spectrum? Is this a single line? Is it broad band? What about it?

(Transparency)

We talk about measuring the spectrum of the light. This is from Bob Hiller's work at UCLA and published in PRL. It shows the spectrum of sonoluminescence as a function of wavelength.

There are two things I want to point out. One is that data point right there. About the same wavelength where we see the decrease in the intensity is where water starts to absorb highly in ultraviolet. So even though the dashed line comes down there, it does not mean there is a peak in the spectrum; that is the cutoff of water. We do not know what the spectrum is doing out there.

This solid line is what you would expect if you treated this system as a blackbody radiator. The temperature of his blackbody was 25,000° -- this also got people's attention. So the first spectral measurements showed that to a large extent you could fit it with a blackbody that had a very high temperature.

(Transparency)

We tried to push the spectral range down a little bit. Quartz cuts off at a little bit lower in wavelength and you can squeeze a few nanometers more out of quartz fiber.

(Transparency)

We measured the spectrum -- this is ours compared to the UCLA one. There is pretty good agreement, considering they used LA water and we used Monterey water.

This is our blackbody fit, 20,000° or so. It basically confirmed what they saw. Then we started varying things. It turns out that Felipe's original solution was a mixture of water and glycerine, so we measured the spectrum as a function of the percentage of glycerine in the water and found that no glycerine was a lot brighter than 25% glycerine and 40% glycerine made a very dim bubble under equivalent drive conditions.

If you put things on a log plot, it looks like something funny happens, the slope on the glycerine seems different from the slope on the water -- do not know why. By the way, this is the answer I am going to give from now on -- "don't know why."

(Laughter)

What else can you do? You can drive the pressure harder. You can put different muck in the water. You can also cool the water down and heat it up. The guys at UCLA looked at what happens when you have 22° water and 10° water.

(Transparency)

Here is the 22° water and here is the 10° water. The spectrum has apparently shifted more toward ultraviolet. Again, there is this cutoff, probably due to water absorption. What else can we do?

(Transparency)

Here is the number of photons per SL flash as a function of the temperature of the water, going from room temperature down to almost freezing. You see, if you look at the intensity at room temperature or number of photons per burst, it is about half-a-million. Cool the water down a little bit, it goes up by an order of magnitude more light. So colder bubbles are brighter and the spectrum has shifted more toward ultraviolet.

So now you have got this little classic system ripping out pulses of light at incredibly synchronous periods and it is probably in the UV, high-energy photons.

DR. MIGLIORI: When you compute the temperature from the spectral density, it is important to correct for all sources of absorption or reflection, so was that done here, because if you have diluted blackbody radiation, you will get the wrong temperature.

DR. ATCHLEY: No, the use of blackbody is just a guide to the eye. There is no reason to believe that either a single temperature would fit it or that it is in equilibrium, necessarily, to be blackbody. It is just what was done because it could be done.

DR. MIGLIORI: Because you could way underestimate the blackbody --

DR. ATCHLEY: Yes, I do not want to lead anybody to believe it is a blackbody spectrum. We will see later that, in fact, it may not be at all.

(Transparency)

You are getting a feel for what is happening in sonoluminescence from about 1992 to 1994. It is cooking. Let us try a dash of salt. Let us try some chocolate chips. What else is in the spice cabinet? What else can we do to it?

One thing you can do is put different gases in the water, dissolve different gases in the water and hope they show up in the bubble. Remember what Professor Garrett said about argon? Professor Bass said the same thing. If you look at the natural concentration of gases in air, you

have got nitrogen at 79 percent, oxygen at 20 percent, about a percent of argon, some CO₂, depending on where you live, etc.

(Transparency)

So let us play around with gas concentrations or gas mixtures. This scale is intensity of sonoluminescence normalized to what you would get from an air bubble. This is the pressure at which you saturate gas inside a liquid, so these are all degassed liquids. One atmosphere would be 760 mm out here somewhere, so these are all degassed liquids.

Let us try starting from the top: Air is mainly nitrogen, right, so take the air out and put nitrogen in, look at sonoluminescence. There it is. There is what you get with real air. Here is what you get in nitrogen: A factor of 50 smaller intensity.

Try nitrogen and argon, 1% argon as you see in air. Boom. Way up there. If a little argon is good, let us put a little more in, let us go to 10%. It comes down a little bit. Let us go with pure argon. It comes down a little bit. If 1% is good, maybe we want to go a little bit lean. Let us go a 10th of a percent. Boom.

It appears that not only was Felipe lucky, he was fortunate, that there happens to be 1% argon in air, because if there had not been, he might not have seen it.

These data show peak the intensity, how many photons come out for different gas mixtures. Let us try measuring the spectrum now with different gases. These next graphs are really horrendous to look at.

(Transparency)

Here is the spectrum of pure argon. Here is the spectrum in 1% argon and nitrogen. Here is the spectrum with 2% xenon -- remember the spice cabinet? Let us just sprinkle stuff in. What is interesting is that the spectrum seems to depend on gas concentration but, more interestingly, remember that peak? Is there a peak or not in the U.V? Look what happens here when you put 1% helium in nitrogen. Lower intensity, but notice that there is a peak, and it occurs in a spectral range where we have some confidence that there really is a peak.

It still does not tell us where the air spectrum peaks, but it does say this thing peaks somewhere, which is nice to know.

DR. BROWN: How do you know the concentrations in the bubble? Just from faith by what you dissolve in the liquid in the transfer function?

DR. ATCHLEY: The question is how do you know the gas concentration in the bubble. What you do is you degas it as well as you can to get as the air out and then you bubble a gas mixture through it at a known pressure head. Then you just hope that Henry's law takes care of it and you get the same stuff inside the bubble.

But this is tricky, because diffusion rates in water are gas dependent. How does argon diffuse across the bubble wall compared to nitrogen or anything else? Just because it is in the water does not mean it is in the bubble.

What else can you do?

(Transparency)

One more knob to tweak. Let us change the water from ordinary water to heavy water, and instead of water vapor let us put deuterium vapor inside the bubble and stand back, get behind your lead sheet. Who knows what is going to happen?

So we measure the spectrum. This is heavy water. This is water water. The solid line is deuterium gas, D_2 , dissolved into the water, and the dotted line is hydrogen, H_2 , dissolved into the water. You measure the spectrum and it is different for heavy water from ordinary water. Why? Do not know.

DR. GARRETT: With solid xenon you had a heavier --you saw a shift in --

DR. ATCHLEY: Yes, but xenon to air is a big change in mass. H_2O to D_2O is not that big of a change. If it were chemistry, a chemist would tell you it does not matter.

DR. GARRETT: That is a common misperception. Physical chemistry does matter on isotopic mass, and particularly hydrogen and deuterium, because all the transport coefficients --

DR. ATCHLEY: That is true. You are right about that.

DR. GARRETT: The standard knee-jerk reaction is there is no chemical effect. That is absolutely incorrect.

DR. ATCHLEY: You are right. Thanks for pointing that out.

So far I have painted a picture of stability. Let us chip away at that. If you start to change parameters a little bit, what you will see is a bright bubble and then it apparently goes away and then, boom, it comes back with no change in drive.

(Transparency)

This graph shows intensity as a function of time. The time scale is six seconds, two to eight seconds. Remember, this is being driven at tens of kilohertz, so a second is a lot of cycles for this bubble. What it says is that for 10,000 cycles everything behaves the same and then all of a sudden something changes, and it stays that way for about another 10,000 cycles, and then it changes, again, but it goes back to where it was and repeats this over and over.

Why?

("Don't know.")

DR. ATCHLEY: You are getting it!

What I have done in this part of the lecture is to give you some indication of the property of the radiation. What I want to do next is see if we can really investigate what the bubble motion is.

The radius-versus-time curves that I have showed before were just computational fantasy. Let us see if we can measure what the bubble really does.

(Transparency)

We started off talking about how you might expect the radius of the bubble to change as a function of time. Then we shifted to what properties does this light have? Clearly, to have any explanation of this phenomenon, or a complete explanation, you have got to couple the dynamics of the bubble to the emission of the light. So let us spend a little bit of time talking about how you might investigate the bubble dynamics.

Remember this type of picture. The sound drive -- somebody asked me this and I should emphasize it even more -- we do use the same field to levitate and to drive. You could do it differently, I guess, but that is what we do. This is the amplitude of the sound where the bubble is and this is the theoretical response, and this is where light comes out.

How could you measure that? A lot of people in this room have spent more time doing this than I have, but if you look at -- for instance, you have this bubble levitated and you shine a laser on it from some angle, you make sure the beam is big enough to cover the whole bubble uniformly, and now you have a detector that you move around as a function of position.

You look at how much light you see as a function of angle. You get some curve that looks like this, where you get maxima and minima all over the place. This is calculated with Mie scattering theory.

What it tells you is that if you have a narrow-angle detector, unless you know the angle very precisely, you do not know what you are supposed to be measuring, so people have taken the approach of using a wide field of view and averaging this curve over a whole range of angles. In that case you get something like the solid line, and that is a number that basically says the amount of light scattered into some large angle is roughly proportional to the square of the radius, that is, the cross-sectional area, and there is some numerical factor that multiplies it, and you get that factor from here.

DR. MARSTON: To be sure, the solid line there is a geometrical optics calculation rather than simply an average of the Mie theory. There are some important differences, but those are subtleties.

DR. ATCHLEY: Thanks, that is a good point.

(Transparency)

If you look with a big lens, you get a nice curve like this. This is also Brad Barber's work, or maybe Bob Hiller's. This line shows what happens when the laser is off -- that is the background noise in the system. You turn your laser on and you see this curve, which is pretty reminiscent of the radius versus time (except it is upside down).

If you do not filter out the sonoluminescence, then you also get a little peak when the light is given off, so it says that the bubble grows and collapses, and somewhere near the collapse light is given off.

What we want to do is see if we can see any change in this gross behavior of the bubble that is correlated to the emission of the light and try to get some feel for how they are connected.

(Transparency)

You can take that scattered optical signal given by the points here and you can fit it to a solution of the Keller equation. To fit it, you have two adjustable parameters, the drive amplitude and the equilibrium size of the bubble. These show that agreement between a solution of the Keller equation and the scattered signals -- actually, I have got to be careful. What they have done is they have taken the square root of the scattered signal to get radius instead of cross-sectional area.

Then they fit it with a solution to the Keller equation and from the best fit they say the drive is 1.50 atm and the bubble is 7 microns. The agreement is pretty good. So this gives some faith, anyway, that you can measure the gross behavior of the bubble.

But of course, what is important is what happens down here near the minimum in radius, but let's not worry about that yet.

(Transparency)

These are graphs of radius versus time inferred from scattering measurements for different drive amplitudes. This is the radius of the bubble and this is the background signal with no bubble there (with maybe the laser off), and this is the intensity of the luminescence.

What they show is that if you drive a bubble low enough you see this RT curve change, the maximum radius gets a little bit bigger, but no light. When you crank it up a hair more, light turns on, the peak radius drops, and it looks like the equilibrium size also drops a lot and the rebounds go away.

So there is a change as you cross over from no luminescence to light, and then you drive it harder, you see changes in the RT curve, the peak radius gets bigger, the equilibrium size then gets a little bit bigger, light grows more and more light. You turn it up enough and, boom, the light goes away. This is for an air bubble.

(Transparency)

If you do it for an argon bubble, pure argon, you do not see that distinct change. You just see a state with no light and then it just sort of turns on gradually, but there is no real obvious change in the dynamics of the bubble. Why? What is the answer?

("I don't know.")

(Laughter)

MR. HALLAJ: The 1.3 atm is not just where you stop plotting. As you said, the bubble goes away. That is like another big "I don't know why," because you go from the brightest sonoluminescence to nothing.

DR. ATCHLEY: Yes. Why? "I don't know."

One disadvantage to the technique they used to get radius as a function of time is that they used the hydrodynamic solution to the equation to fit the radius. It would be nice if there were an independent technique of determining the radius of the bubble independent of any other kind of technique. I will not get into details; we worked on that at the Postgraduate School -- there are some plots like that in your notes.

(Transparency)

We will skip over those, but I will say the most important thing about that work was it got us on the cover of Applied Optics. So there is the laser coming through, there is a bubble there, good PR. Let us skip over that stuff.

Remember I said the important thing is what happens at that minimum, when that bubble is crashing down, what happens? That is where you want to start looking. The rest of it is not too interesting.

(Transparency)

If you take their scattered measurements (this is, again, Putterman's group), this is the RT curve. You blow up this box here and you get this graph and now you blow up this box and you get this graph. You would like to get resolution down to here.

This little bit is an interesting artifact of the measurement system. I will not get into it, but it shows that if you really want to understand the system, you have got to understand the measurements you are making. They actually said, shown from this, that the luminescence from that bump is a little indication that sonoluminescence occurs maybe 10 ns before the minimum bubble radius. It turns out that is an artifact of the measurement technique. They have not retracted that in print; but they did at an ASA meeting. But that is an artifact of the system that we can talk about in the evening some time.

Remember that I said that sonoluminescence from transient cavitation, random stuff, has been known for a long time. Why don't we compare the two? That is, how does old-fashioned sonoluminescence compare to Gaitan-style sonoluminescence?

(Transparency)

Tom Matula and others at the University of Washington did that comparison. Ken Suslick is a sono-chemist and he has been doing luminescence for a long time and has been measuring spectra for a long time.

(Transparency)

When he measured spectra from multiple bubble random transient sonoluminescence, he would see well-defined emission lines. You put salt in the water, You see a sodium line. You put potassium in the water, you get potassium lines. You get these hydroxyl free radical lines. You see these peaks -- always there. You can use this to determine all sorts of stuff about what is going on inside the bubble.

He saw the single bubble SL spectra and said, "It's got to be garbage." He said, "There have got to be lines." He got together with guys at U.W. and he wanted to show that their other spectra were wrong, there had to be lines. So he uses his fancy system and measures this spectrum from single-bubble sonoluminescence. Not even a hint of lines in it.

Why?

("I don't know.")

DR. ATCHLEY: Actually, there were details about the system we can talk about maybe in the evening, there are reasons why there might be differences.

(Transparency)

What else can we tweak? Go to the CRC and look up all the physical constants and let us see which ones we can tweak. One of them is gravity. This is Tom Matula, a former PASS student, in the KC135 doing the parabolic arch -- what do you call it? the vomit comet? And in here is a sonoluminescence apparatus, I think.

What they did is took the bubble up in this plane and measured the intensity of light as a function of the apparent gravitational force on the bubble. Why? "I don't know."

(Laughter)

(Transparency)

This is some of their work. Tom was nice enough to send me this stuff, even though it is preliminary and I am supposed to say that this is out of the goodness of their hearts and you should not believe anything too strictly.

This is kind of interesting. This is the gravitational force in that plane as a function of time. Here is the intensity of the sonoluminescence as a function of time. You can see it is not an Albert Migliori measurement, but there is some oscillation here.

What does it show? It is interesting that there does seem to be some effect. Remember, when we did buoyant levitation, there was buoyancy in there. When you turn the buoyancy force on or off, the equilibrium position of the bubble changes up and down. So maybe that has got something to do with it. Who knows? But I thought you would enjoy seeing that, and I can tell Tom that he has been presented to the meeting.

(Transparency)

Let us drive it a little bit off resonance and see what happens. This is some stuff we did. Do you know about period doubling, chaos, and that kind of stuff? If you have a harmonic

oscillator and it is linear and you drive it and you measure the position at a certain phase in the drive, if it is linear, it is always in the same place? It is called period one motion. If you drive it harder, it becomes nonlinear, maybe the oscillations get a little bit strange, and it takes two cycles to come back to where it started? It is called period two motion; you can do it, again, in period four motion. There is this period doubling route to chaos.

What we did is get a bubble going, make it stable, and then detune the resonator slightly. What we found is this. Under certain drive conditions, the arrival time is always there, flash after flash after flash.

If you detune the resonator, then you can get into a motion where the arrival time is not the same all the time, but it repeats. It goes from here to there -- here to there -- here to there. It wobbles back and forth: period two motion. But there are only two well-defined values -- here or there.

You detune a little bit more and you go from period two motion to period four motion. Then you get this sort of smeary quasi-periodic motion. If you detune it enough, then you get to chaos. These graphs are phased plots -- it is kind of hard to see, but period one motion shows one dot, period two motion shows two dots, four shows four dots. Quasi-periodic motion shows this sort of closed-loop path. Then this phased plot degenerates into this chaotic behavior.

What is interesting about this? Lauterborn has done a lot of work with chaotic behavior of bubble clouds, multiple bubble systems, but there has never been an observation, as far as I know, of period doubling or chaos in a single-bubble system. We did not do is go back and try to relate it to the motion of the bubble, but it does point out something that a theory would have to take into account.

That is enough experimental stuff. What do we think is going on? Or what do some people think what is going on?

(Transparency)

If you go back to the RT curves, there are several ways to model this system. One is to use a Keller-like equation to talk about the motion of the bubble wall and then try to couple that motion into the gas. You treat the gas in some sophisticated way or not, but that is one way you can do it.

This is work from Wu and Roberts at UCLA. They used a hydrodynamic equation to drive a bubble, the wall of the bubble, and then couple that bubble motion into the gas dynamics.

(Transparency)

This graph shows the pressure inside the bubble as a function of radius -- there is the center of the bubble, this is a micron away from the center, at different times. These are pressure profiles in the bubble as the bubble shrinks. You can tell it is shrinking, because these lines start at smaller and smaller radii.

This shows that at some point you get a discontinuity in the slope, which means that there is a shock wave being formed. This bubble collapses and starts to pile up gas in front of the bubble wall kind of like a snow plow, and eventually the gas shocks and propagates inward.

(Transparency)

This is the radius of the bubble, the bubble wall motion as a function of time, and this is the radius of the shock front as a function of time. Their calculations show that the shock gets ahead of the bubble wall, and shrinks to a very tiny size.

What they said is that if you have this high a pressure and density and temperature and you assume the stuff inside is ionized, then you use some standard equation from plasma physics to get how much light should come out.

Let us take it one step further. Let us still model the motion of the bubble wall, but let us start from a numeric point of view. I have got this piece of glass. I have got water inside. I have got a bubble inside that. I know what water is, I know its equation state. I put gas in there. I try to use the fanciest equation state I know. I am going to drive the wall and just let the system go numerically, and let us see what happens.

(Transparency)

This is work from Moss at Livermore. I have put a bunch of graphs in your notes. I should point out these are all preliminary, but they are indicative of the kinds of things he is working on.

Instead of driving steady state, he starts from rest and drives this bubble, and this shows the radius of the bubble as a function of time, where he starts the sound out from the glass wall. It propagates to the bubble. The bubble radius grows and collapses, kind of like Rayleigh-Plesset.

This is a graph of the radius of the bubble, the acceleration of the bubble, and the velocity of the bubble wall as a function of time. The acceleration of the bubble wall is going to become important later -- and you can read the scales and figure out what the numbers are.

(Transparency)

He shows the pressure and the density inside the bubble as a function of time.

MR. HALLEJ: How is he calculating those? Is he using the model or a direct numerical simulation?

DR. ATCHLEY: It is direct numerical simulation and he is using equations of state. This guy is a bomb code calculator with nothing to do, so he is using all these fancy computational techniques and equations of states of things at high density, temperatures, and pressures, and trying to do the best he can.

(Transparency)

I cannot explain all these details, even if we had time, but he tries to take into account ion temperature, electron temperature, and radiation temperature inside these equations of state. This

gas gets hot, it ionizes, nasty stuff happens, but he is doing the best he can to put it all together, but not using hydrodynamics to couple the wall. He is setting up the system and letting it go numerically and seeing what he gets, and this is what he claims is the most honest calculation he can do.

With this model you can get extremely high temperatures and pressures and you can get light somehow. Maybe it is not too surprising. I guess I would say that if you had to poll people, they would probably say this shock-wave model is about the best game in town right now.

But I will still point out that people have not predicted anything from it yet; there has been no prediction of a spectrum that has been directly measured.

DR. GARRETT: No temperature-dependence gas?

DR. ATCHLEY: No, there are no predictions from this stuff, but it seems reasonable.

By the way, these temperature scales are in electron volts, and an electron volt is about 10^4° , something like that, roughly, so when he says 20 eV, he is talking about $200,000^\circ$ -- at high pressure and high density. What comes to mind now?

(Transparency)

Well, the "F word." Fusion. About this time all sorts of stuff hits the newspapers. Can you get fusion from these devices? Seth put in a paper a few years ago that if you have these high temperatures and pressures you might get a neutron per hour, a neutron per day.

People try to tie it to cold fusion, but if there is really fusion to come out of these things, it is cold: This is honest-to-goodness hot fusion.

(Transparency)

There is this credible magazine with a story title, "pursuing the cold fusion genie", like something out of chain reaction, all sorts of great titles and articles --

(Transparency)

-- "Tapping the Light Fantastic" --

(Transparency)

"Science News," "Inferno in a Bubble," all these great things.

Willie Moss had this idea that under these conditions I just showed you maybe there was not enough to get fusion out. He said, "Let's kick the bubble." His favorite term is "Give a child a hammer and everything becomes a nail."

(Transparency)

So he took this bubble -- this is the acoustic drive amplitude as a function of time -- and asked what if you give it a high spike, a kick, at just the right time. What happens then?

(Transparency)

This is a radius-versus-time curve. It looks like the one I showed you before, except when the spike hits, this collapse is much more violent and much more rapid and compresses things a lot hotter.

DR. MAYNARD: What is the minimum of the bubble radius, ballpark?

DR. ATCHLEY: What does he calculate?

DR. MAYNARD: His best guess.

DR. ATCHLEY: Seth thought it might be half a micron. Seth put in van der Waals equation of state and the hard-sphere limit is what keeps the bubble from collapsing. Is it still a bubble, then? I do not know. It all gets tricky. When things are happening this fast on such a small scale and under such extreme conditions, is it even a bubble? Does it even make sense to talk about bubbles?

DR. MIGLIORI: What about polarization of the water molecules on the surface of the bubble? This is well known, you know, shock polarization, it is a big effect.

DR. ATCHLEY: People have suggested that.

DR. MIGLIORI: That produces a huge charge on the inside surface.

DR. ATCHLEY: The real kicker on this stuff is the gas-concentration stuff. Why should putting 1% argon inside of the bubble muck around with that too much? And why is it 1%.

DR. MAYNARD: Has anybody tried some liquid other than water?

DR. ATCHLEY: Yes, there was something -- gee, I always forget. On Gary Williams' home page they say they found sonoluminescence from a xenon bubble in some alcohol at -115°C .

DR. MARSTON: Actually, Seth had a paper in either The Journal of Chemical Physics or The Journal of Physical Chemistry last -- roughly November or December.

DR. ATCHLEY: Do you have the reference on that? I was trying to find it and I could not.

DR. MARSTON: I have a reprint in my office. Anyway, it is November or December of 1995 in either The Journal of Physical Chemistry or The Journal of Chemical Physics showing stable single-bubble sonoluminescence in a number of liquids, alcohols, nonaqueous --

DR. ATCHLEY: But it is at very low mean partial pressures, right?

DR. MARSTON: Yes.

DR. ATCHLEY: All right, so the point of all this is you can apply fancy computational stuff and apparently this thing could give off neutrons. You are not going to cook a chicken, but who knows? People have looked for neutrons and they have not seen any.

DR. MARSTON: In regard to that, Willie Moss had a paper in Physics Letters this spring on more theory calculations behind that effect -- I do not have that reference.

DR. CHAMBERS: Has anybody ever done the experiment with the extra spike in the sine wave?

DR. ATCHLEY: No. They were trying at Livermore --at least that I know of -- but the experiment is lagging.

DR. BROWN: But you said earlier that you get spikes because you have got neutrons coming in and that sort of sets off transient cavitation. There are all kinds of spikes coming from cosmic rays, right?

DR. ATCHLEY: Yes, but this is sort of a focus-concentrated spike in the water that is concentrated on this center.

DR. BROWN: Some effort.

DR. ATCHLEY: One would presume, yes. I mean, you can think of other ways to kick it, too. But it is an interesting concept and if you believe this computational stuff, it is very interesting.

All right, I want to spend a few minutes talking about what other people think may be going on. Hank talked about three things: stuff he knew about, stuff he sort of knew about, and stuff he did not know about. I have talked about one thing: stuff I do not know anything about.

Let me tell you about three or four of the things that people think may be going on. I was giving a similar talk at the University of Texas, while I was on sabbatical, to the physics department, and this guy raised his hand -- it was toward the end of the hour -- and he said, "I know what's going on, but I have to leave, come find me," and he left. This guy was Lothar Frommhold and he is an expert in something called collision-induced emission.

This model is pretty interesting. It has also got some problems. But here is the picture. Under ordinary conditions, take an atom like argon. You have an electron cloud here and you have got the nucleus and everything symmetric, and you have no net dipole moment of this atom. Dipoles are what radiate in electricity and magnetism.

Suppose you collide two argon atoms, what happens? You start off and they are nice and symmetric. As they collide, what happens to the electron clouds? They see a like charge and they repel one another; they distort. So during the collision the cloud may look more like this, right? It gets distorted from the collision.

Now you have got a dipole moment, and that can radiate. Moreover, the dipole moment changes on the time scale with the atomic collision, so you have no dipole, you collide, you get a dipole, get radiation after they collide, it goes back to where it started and no radiation, so it is quick. Interesting.

The more interesting thing is suppose I do this with two similar atoms. Clearly, I induce a dipole moment here, but because it is the same atom, I induce the opposite dipole moment on the other one. So the dipole moments are the same, and they are close to one another, so they essentially cancel; they can make only a quadrupole contribution. That is going to be a much less efficient optical radiator.

If you have similar gases that collide, this process will not amount to much. But if I collide argon with nitrogen, then the dipole moments do not cancel as much and now you can get more radiation. So this mechanism is sensitive to the concentration of the types of gases, which I think is the most interesting part of it.

The down side is that if the gas is -- this requires relatively cool gases, gases that do not ionize. As soon as you let electrons loose, their dipole moments dominate and that should blow this mechanism out of the water. If you believe any of those shock-wave stuff, you would expect the temperatures to be high enough to ionize. But it is interesting. He is still working on it.

DR. MARSTON: Anthony, do you want to say anything about applications of CIE to stellar atmospheres?

DR. ATCHLEY: That is Frommhold's field, not mine.

(Transparency)

Moving on, here is an interesting title: "Sonoluminescence and the Heimlich Effect." Do you know what the Heimlich effect is, or the Heimlich maneuver? When you get stuff caught in your throat and -- oooh -- it comes out? It is the same Heimlich, and they referenced The American Journal of Medicine, or whatever. I actually have not read this whole paper, but it is pretty interesting.

(Transparency)

Two or three days ago Prosperetti sent some stuff. He says that here is what is going on. There is a big problem with stability. You get this big imploding spherical shock wave, it is not going to stay spherical, and this is what inertial confinement fusion has been fighting for years -- with billions of dollars. They cannot get their shock waves to stay symmetric.

So why would you expect a simple little bubble to be the magic solution? This picture shows this bubble grows and collapses and as it collapses you get this little distortion. From here to here you get a bump that forms. You watch this bump that eventually becomes a jet. This is the same jet that Rayleigh found in 1917 that eroded propeller blades. So the jet has to come back at the end of the lecture.

This jet goes through the bubble wall into the gas and impinges on the fluid on the back side. It is so fast and the emissions are so extreme that it ruptures the liquid, and it is actually a physical fracture of the liquid that causes the radiation. It is brittle under these conditions and time scales and it is ripping these molecular bonds that cause the radiation.

DR. BROWN: I have a question. If it was that asymmetric, then it would form two bubbles when it came back, right? It would not come back to a spherical resonator to do the same thing.

DR. ATCHLEY: It is like a donut. Who really knows? What is interesting about that? I mean, you bring up a good point. If this collapse is so violent, if the bubble does distort itself and shatter, you might expect that when it comes back together maybe some of the pieces are missing.

So why should it behave the same the next cycle and 10,000 cycles later? You know, it has got this long-term memory. Whatever happens down there, it is not enough to affect this long-term memory of what it is doing. It is really interesting, especially this chaotic stuff, this period two motion and period four motion. It shows that somehow, whatever happens down there, the system knows where it came from and can repeat it, again.

DR. BASS: I have not read this paper as carefully as you have, but does he carry the calculation through beyond -- that gets through in this simulation?

DR. ATCHLEY: Into the radiation?

DR. BASS: No, beyond to the next cycle.

DR. ATCHLEY: I do not really know. Oh, yes, you know, Ani, you speak up for Andreas.

MR. WORLIKAR:you see the jet formation is the last stream of that picture, so when you go through there is no jet formation until the last stream....

DR. BASS: Why don't they do the computer cycle when there is an oscillation?

MR. WORLIKAR: Because you do not know what happens after this jet forms.

DR. HARGROVE: Lauterborn's movie showed that this thing can go on....many cycles, so it does not have to break up.

DR. ATCHLEY: Right. Those drive conditions are a lot different, though. Lauterborn is not looking at little bitty bubbles.

All right, here we are, near the end of the lecture, near the end of the century, near the end of the millennium.

(Laughter)

I have taken a two-and-a-half hour lecture and made it seem like a hundred years long. Where were we 100 years ago in physics? I know, you were not there, but there were things we did not understand -- maybe one thing we did not understand. We thought it was over for us. But there is this little problem of optical radiation and atomic structure of material that got us through the next 100 years.

I have tried to give you an indication of classical approaches to what could cause sonoluminescence, but there are people out there who are dying to get an \hbar into the picture. Julian Schwinger, at the time of his death, was working on sonoluminescence, and he had a different picture of it.

Similar work has continued, and recently there was a series of papers, one in Physical Review Letters, one in Physical Review and other places, asking is sonoluminescence a

manifestation of quantum vacuum fluctuations? That is a mouthful. Let me explain things from a very simplistic point of view.

People are probably familiar with quantization of energy levels in systems, where the energy is proportional to the frequency of oscillation. Depending on what mode you are in, the energy is higher and higher, like standing waves. There is $n \hbar \omega$, but there is also a factor of a half $\hbar \omega$, where ω is the optical frequency, that hangs around. So even when you are in the ground state of absolute zero, there is still some energy; it is called the zero-point energy.

If you take a state -- this is in the ground system at zero degrees and, by the way, I am not a quantum field theorist and anything I am about to say now is even more out on the limb than it was with bubble dynamics -- if you take a state where there are no photons emitted, then you would think that there would be no electric field existing.

And if you ask what the average value, the expectation value, of an electric field is in the ground state, it is zero. Fine. If you ask what the fluctuation is, or the energy -- that is, not look at the average value of E , electric field, but look at E^2 . It turns out the answer is not zero. There can be huge fluctuations in this vacuum where there is nothing. There can be huge fluctuations of electric field.

One manifestation of this is something called the Casimir effect, and we can spend time, maybe tonight, talking about it, but basically it says if I take two conducting plates, uncharged, and get them close enough together, even at zero degrees there will be a force of attraction between them, and this force of attraction is a manifestation of fluctuations in this vacuum state.

There are other connections to zero-point energy. Black holes, for instance. Black holes, people think they suck everything up. They actually can emit radiation as well and that radiation mechanism is also connected to vacuum fluctuations. Radiation from black holes depends on the value of the gravitational acceleration at the black hole, which can be extremely huge, extremely high.

It turns out that if the acceleration is on the order of 10^{20} g's, you can get this spectrum radiated from a black hole, and I do not know what the temperature is, but it is some measurable temperature. So what does that have to do with bubbles?

Schwinger's idea was, if I have a bubble in a fluid that has a volume like this, it has got some certain energy. If I have a bubble in a fluid that is this size, it also has energy, but it is a different energy. If I go from here to here, the energy changes; that energy is radiated away from the photons. That was his point of view. The photons came from this change in the vacuum energy.

This person has taken it a step further and tried to couple this mechanism to the dynamics of the bubble wall. If you look at the accelerations the bubble wall experiences when it turns

around, they can be pretty big. Frankly, she says they are bigger than what Moss' calculations say, by orders of magnitude, but let us just continue the thought.

There is a big acceleration and this acceleration can lead to radiation. It turns out, if you look at her numbers, with huge accelerations you can get a spectrum that looks kind of like a blackbody with a peak temperature of about 40,000° under extremely radical acceleration conditions.

So maybe you say, well, we cannot have those high accelerations. Maybe it is not the bubble wall that does it, but it is maybe there is a shock, and when it collapses, there might be a big acceleration in this rebound. It also depends on the difference in index of refraction between water and, say, air inside the bubble (water has an index refraction of 1.3, air is nominally 1, so there is a 0.3 difference in index). You have to have a difference in index of refraction in order to have this radiation mechanism.

Or maybe inside the shock, where you have real high densities and pressures. Maybe across the shock front there is a much bigger index of refraction change than the bubble wall change, which could reduce the amount of acceleration you need to get the same radiation. It is like a binary counter; if you poll people, it is either the best thing in the world or it is complete garbage.

I cannot tell you the answer. I think it is interesting and I think it is maybe more interesting from an historical perspective at looking at where we are now, where we are going, and think that maybe -- maybe -- sonoluminescence is the first direct manifestation of quantum background fluctuations. Maybe.

(Transparency)

Let me show you two more graphs and we will be done. John Kordomenos, a graduate student at the University of Mississippi, is doing an interesting experiment, which is also preliminary. He is looking into the microwave region of emission. He drives the levitator with a horn and he has got a little dipole antenna stuck in this flask and he is looking for radiation in the microwave region.

(Transparency)

They have done a calculation, and this is also preliminary, but it is being refined all the time, he and Bruce Denardo and other people, and they extend the spectral calculation for different mechanisms out into the low-frequency range, microwave regions.

Here is bremsstrahlung. Now, what is bremsstrahlung? If you believe that you ionize a gas inside and have a bunch of free electrons floating around, that the radiation mechanism is the acceleration of these charges when they collide, basically, that is the mechanism called bremsstrahlung.

This is what you would expect this mechanism, the spectral behavior of it, to be at microwave regions. They plotted three other mechanisms: blackbody, this Casimir stuff I just talked about, and just to make me happy they put in collision-induced emission.

They measured how much emission you get from microwaves. You people from Mississippi correct me, if I am getting this wrong, but they did not see emission in the microwave region. But they have compared the noise level of their system and it is somewhere around here, which means if you believe their calculations, it would tend to rule out bremsstrahlung as a radiation mechanism. They cannot say anything about the other mechanisms.

This would be pretty significant, and what this would tell you, then, is that it is probably not this ionization and plasma stuff that is doing it. Who knows? But it is interesting. Do you want to correct this?

DR. CHAMBERS: No, but I just want to make sure that I understand it. If I remember correctly, they adjusted the numbers of the other models so they all fit what they see in the visible light spectrum or what is actually observed. They all look very visible, because they have to.

DR. ATCHLEY: Yes, this, again, is preliminary stuff. The stuff I have told you in the last half-hour is just what is out there in the literature, but it is interesting. If you can start picking off these mechanisms one by one -- what did Sherlock Holmes say? -- if you exclude the impossible, what you are left with however improbable is the answer.

MR. HALLEJ: So if you can exclude the bremsstrahlung for sure, what does that imply now for the shock-wave theories?

DR. ATCHLEY: I think it would be trouble.

MR. HALLEJ: So it is bad news for that?

DR. ATCHLEY: I think so. But all these models have troubles, and this is preliminary. This is just where we are today, June 25, 1996. And with that, I will end. Questions?

MR. DOUST: This might be sort of naive, but if you are emitting this light energy from this, where are you bringing energy back in to make it equilibrium?

DR. ATCHLEY: Presumably from the acoustic field.

MR. DOUST: In some of that it sounds like you are saying you are generating more light energy than -- in terms of energy, the more light that goes out in that one three-dimensional plot you are not always getting a relationship between the acoustics and the light.

DR. ATCHLEY: I cannot answer your question satisfactorily, but there is one thing I did not really stress enough that people stressed early on in this field. If you take the acoustic intensity and you evaluate it on a per-molecule basis, acoustics is a pretty weak mechanical phenomenon.

If you spread it over 10^{23} water molecules, gas molecules, you do not have a lot of energy per molecule from acoustics. If you say that in order to emit light, the energy of the molecule has to be an electron volt or so, typical energies of photons, and if you look at the ratio of

that optical output per molecule to the acoustic input per molecule, the ratio is something like 10^{12} ; there is a huge concentration of energy from this mechanically weak sound field into the molecules. I do not think anybody understands that.

DR. MAYNARD: If you tried to shoot down your own theory by having quadripoles instead of dipoles....molecules, those things are moving relative to one another by -- in fact, what are those velocities? Just guess.

DR. ATCHLEY: I would not guess. But estimating collision times, what are they going to be? Nanoseconds, these high temperatures and pressures? I do not know. I do not know those numbers because I just would not believe them if I calculated them.

I shot it down for the same reason I shot down the rest of them; it does not predict anything yet. Part of the trouble is, in order to do these calculations you have got to know a lot about the properties of atoms and molecules under extreme conditions and we just do not know these things.

DR. MAYNARD: He tried to shoot down his....collision theory, because when you have identical atoms there is an equal and opposite dipole moment induced in each one, but they are moving relative to one another, so if you go to the resting of one or the other, then there is an asymmetry, you do get a dipole moment.

DR. MIGLIORI: There is also a dipole layer. You can accelerate stuff right through it, just from the acceleration-induced --

DR. ATCHLEY: It may still work. I just do not have any reason to like it preferentially just because I happen to be the co-author.

DR. MAYNARD: For a long time people have done molecular beam experiments of all sorts.

STUDENT: (Inaudible.)

DR. ATCHLEY: That is also a big problem with stability of bubbles that are even 5 microns in size at gas saturation pressures near atmospheric. If you go to low saturation pressures, like millimeters of mercury instead of hundreds of millimeters of mercury, then the diffusion apparently can account for the long-term stability of the bubble, but that is a big unknown. How do you determine anything like this theoretically?

DR. BROWN: Can you run sonoluminescence for days, weeks, months?

DR. ATCHLEY: The liquid tends to age and the bubbles tend to wimp out after a while. I have kept it for hours and then I wanted to go home. If you sat there and did a marathon, I do not know how long you could keep a bubble. Probably not for more than a few hours, just because if the system is open to the atmosphere and gas can diffuse into it and dirt can agglomerate on the bubble surface -- but I guess if you were careful you could keep it going.

DR. MARSTON: If you put up the transparency with Eberlein's paper on it, just a comment.

(Transparency)

On the right-hand side there is a hint as to why some of the theorists are as interested as they are. If you talk with people who specialize in the quantum theory of radiation, they have had an interest in what is known as the Unruh effect, listed here with a reference. What that was, was Unruh's analysis, if I understand properly.

The radiation from an accelerating mirror -- I am not sure whether it was Eberlein or not, but I think she may have also extended that to the accelerating dielectric. The point is, if you continue on down in the paragraph, she indicates that sonoluminescence may be the first identifiable manifestation of quantum background radiation.

In other words, the Unruh effect appears to be so far unobservable, so this may be the first platform on which that effect would be observed. For whatever reason, for people who specialize in quantum radiation, they had already known about the Unruh effect and it was a curiosity as to whether it could exist.

DR. ATCHLEY: As I understand, which is not worth a whole lot, the Unruh effect, talking about these fluctuations in the electric field, if you look at the spectrum of these fluctuations and you plot out the spectrum (this is intensity in some Wright units versus frequency), the spectrum of this background fluctuation has to go as omega cubed in order to satisfy relativity-type stuff where there is no preferred reference frame.

This, again, is unobservable. This is the spectrum of virtual photons and I guess this all occurs at 0° K, too. If you have a uniform velocity, then you will not see any of this spectrum, because that would mean there would be some preferred coordinate system, and that violates relativity.

But if you have an accelerating coordinate system, then you can get an isotropic radiation of real photons whose spectral shape is something like this. These are real photons and look like a blackbody.

(Transparency)

The spectrum depends on \hbar and acceleration and the speed of light -- when it is that shape and that peak temperature, which we infer from some fit, depends on the acceleration and that, I think, is Unruh effect. So even if you have this uncharged mirror, if it accelerates, at huge accelerations it emits light -- real photons, not just background fluctuations.

If you fit a peak temperature to this, it is something like $\hbar a$, where a is the acceleration, k_B Boltzmann. If a is equal to one g, then this peak temperature is like 10^{-20} K, so if you have 10^{20} g's, then you get a degree out of it. It will not cook a chicken, but it is interesting. Thanks, Phil..

STUDENT: Are we going to see the demonstration?

DR. ATCHLEY: Yes. Let me try to set it up.

FUNDAMENTALS AND APPLICATIONS OF NONLINEAR ACOUSTICS

MARK HAMILTON
UNIVERSITY OF TEXAS

DR. HAMILTON: A number of the previous speakers have mentioned various ways in which nonlinearity affects either their analyses or their experiments. The intent of this lecture is to focus on those nonlinear physical principles that pertain specifically to the propagation, distortion, and interaction of sound waves.

(Transparency)

The approach I will be taking is to use primarily governing equations that are derivable directly from the classical equations of fluid mechanics, and then to illustrate the results of these equations with experiments, oftentimes that were designed specifically for this purpose.

Although a number of these physical principles and equations are disarmingly simple in appearance, frequently their derivations are not, and we will not go into those here. For the derivations of the classical results for progressive plane waves, I refer you to the excellent review article, the *AIP Handbook* chapter by David Blackstock that was sent to you prior to this meeting.

For a number of the derivations for more complex wavefields, for example, diffracting sound beams, I can refer you to the list of references in the second review article that you were handed.

(Transparency)

A good place to start is, in fact, with the physical description of the mechanisms behind the distortion of progressive plane waves. For this I will basically refresh your memory of the description Robert Keolian gave you in the demonstration session.

For the purpose of discussion and for simplicity, suppose we have an ideal gas and suppose we radiate a portion of a waveform like this into that gas. In those parts of the wave that are compressed, the positive parts of the wave, if the gas is adiabatic, the temperature will increase slightly and, therefore, the wave sees locally a slightly higher sound speed.

In those portions of the wave where the gas is rarefied, the temperature drops slightly and that portion of the wave will see a slightly slower sound speed. The zero crossings of the wave are at the ambient temperature and those will propagate at the same speed.

So the net effect of this nonlinearity due to the equation of state is to cause these waveforms - here would be plotted versus time -- would be to cause the compressive parts of the wave to advance forward in time and the rarefied portions to lag behind.

There is a second effect, often equally as important, for progressive plane waves. This is referred to as convection. For a progressive plane wave, the particle motion has the same sine as the excess pressure in that portion of the wave.

In the positive portions the fluid motion is going forward; in the negative portions the fluid motion is going behind. The physicists of the 19th century referred to this as the "local acoustic wind" and described the distortion process due to this as "convection."

It happens that the convection causes distortion in the same direction as the nonlinearity. Again, we get steepening of the positive portions and a falling back of the negative portions. There are a few exotic liquids where these two effects can cancel and, in fact, where the equation of state can overcompensate for the convection and give you a reverse distortion, but these are by far the exception and not the rule.

As this wave propagates, eventually we encounter very steep gradients at the points of the zero crossings in this example. At those points, the effect of viscosity becomes very strong. Heat conduction is enhanced. These losses act locally at the shock front in such a way as to offset the tendency of this wave to steepen forward.

When the nonlinearity is sufficiently strong to maintain this steep slope in the waveform, we call that a shock front. As the wave propagates at sufficiently high amplitude, additional portions of this part of the wave catch up with the shock front. As these portions catch up, this wave reduces in amplitude and acquires a sawtooth shape, much as you saw in the waveguide experiment that Dr. Keolian showed you on Sunday.

Eventually the amplitude of this waveform becomes very weak. The effects of viscosity and heat conduction eventually win out. The rise time of the shock increases. The shock relaxes or, in the frequency domain, you could think of this as a low-pass filtering of the wave. The omega squared thermoviscous dependence of attenuation that Professor Bass talked to you about effectively strips away these high-frequency components that form the shock front and this ultimately low-pass filtered wave, if you will, can return to something more resembling the sinusoidal waveform that was originally radiated.

(Transparency)

How do we describe this process mathematically? We begin with the exact equations of motion for a plane wave. I have chosen an ideal isentropic gas. The equation of continuity is the first equation. Momentum conservation is the second equation. Because we have only one-dimensional motion assumed in this gas, we have only three variables in this case: the pressure, density, and the particle velocity.

With these three equations we can combine them through some rather clever substitutions, argue how to integrate these results, and ultimately obtain a single equation in terms of a single field variable, that being the particle velocity, U .

This equation here is an exact reformulation of these three nonlinear equations of motion. It is quite remarkable that this result can be distilled from those equations. This equation was first derived by Poisson in 1808. In this equation C_0 is the isentropic sound speed and β is the label

we give to the coefficient $\gamma + 1/2$. γ is the adiabatic gas constant and β is referred to as the coefficient of nonlinearity.

If we set β equal to zero in this equation, you will see that the nonlinear term here vanishes. We are left with these two terms. These describe progressive plane waves in the positive x direction.

Poisson recognized that the system can be inverted and that an exact analytic, although implicit, solution of this equation is given by this function here. Note that U is an arbitrary function of itself and it can be rather difficult to invert this equation to get an explicit solution.

Nevertheless, Poisson was the first to recognize this solution, and from this we get our first really fundamental law for the propagation of finite amplitude plane waves: Namely, suppose we take a look at the waveform and put a little dot on one portion of the wave, that is, a part with constant velocity U , constant amplitude.

We want to track that point of constant amplitude. Setting U equal to constant, you can see almost by inspection that the propagation speed of that -- we will call it a wavelet -- is what is given by the denominator here, $C_0 + \beta U$. If β were zero, C_0 would be your phase speed.

What we have here is the finite amplitude propagation law that accounts for the distortion in the previous slide. Positive parts of the wave have higher particle velocities and propagate faster and those with negative ones fall behind.

Moreover, we can check our physical intuition at this point. Suppose, in fact, as Poisson thought -- Poisson actually worked this out for an isothermal gas. In an isothermal gas, pressure is proportional to density, the adiabatic exponent is unity. If I set γ to 1 in the nonlinearity coefficient, β equals 1, what I find is, then, the propagation speed is just a small-signal sound speed corrected by the particle velocity. There is our convection effect in the absence of the nonlinearity in the equation of state.

(Transparency)

This little cartoon here, where it is supposed the louder a sound wave is, the faster it goes, would not have been a laughing matter in the 19th century. In fact, there was an extended series of exchanges in the British journals between the physicists on how to interpret these laws.

The earlier physicists took a look at this propagation law, saw the particle velocity in it, and thought, AHa!, the louder the wave, the faster it goes, but, they said, "That works against our intuition and in our experiments that doesn't happen, our results must be wrong."

In fact, prior to Poisson, both Euler and Lagrange had derived slightly incorrect versions of these results, but they got the main effects there, and immediately concluded a propagation law like we derived and then threw it away, considering that the result had to be absurd.

It was not until, really, the middle of the 19th century when Stokes settled the debate. He published the first descriptions of waveform steepening as we saw in the first slide and was the

first one to identify the fact that a shock wave must be formed and that dissipation must be the competing mechanism to stabilize that shock front. This was not a subtle point.

DR. ATCHLEY: If you go back to your wave equation, you start with the same equations up there, what is it about the one-dimensional propagation that gives you that?

DR. HAMILTON: There is a compound wave result that first follows from these equations - Rehman derived this. To simplify these equations you have to introduce auxiliary variables that are difficult to assign physical significance to.

You have essentially two first-order equations, which, if you combine them, would give you your second-order wave equation. But this simplified form can be obtained only once you assume wave motion in a single direction. You consider this a first integral of the equations of motion.

DR. ATCHLEY: The other equation is where you built linearity into it, it is already linear?

DR. HAMILTON: Yes, it is already linear.

DR. GARRETT: Another way to look at this is you have got a first-order equation here, but it is only for the wave going to the right. You could write another equation, with a minus sign in front of the parentheses, which is the wave going to the left. Take $A + B \times A - B$, you get $A_- - B_-$, and you are back to the wave equation.

DR. HAMILTON: This cuts to a fundamental point here. Superposition does not hold. When you decouple these wave equations, they are radically different from the equations describing compound wave flow. Really, to this date nobody has extracted analytic results for even the simple case of standing waves. These are very difficult problems even for a simple standing wave in a tube. That is a good point.

(Transparency)

Let me show you a slide I made up last night, because it motivates the next one -- you will not find this in your handouts. I want to give you kind of a bird's eye view of, really, the first law of how to solve equations in nonlinear acoustics when no exact solutions are available. We call this the method of perturbation.

The idea is we separate our wave equation with a linear term, on the left, and, for example, group all the nonlinear terms on the right. For the equation I just showed you, this would be our linear wave operator, F , and G is obtained accordingly.

What we do, then, is we expand our dependent variable in an expansion where it is assumed that each term is a very small correction to the one before it. The normal ordering parameter for these expansions in nonlinear acoustics is the acoustic Mach number, typically a source amplitude, velocity amplitude, over the sound speed.

To give you an idea of the magnitude of this parameter and the assumption of its being much less than 1, a value of 0.1, of 1/10, for the Mach number is equivalent to about 174 dB in air. For

very small Mach numbers, from a mathematical point of view, you can get very substantial nonlinear effects. So these expansions are often very helpful.

When you substitute this type of expansion into this equation, you find that at lowest order you are simply solving a linear wave equation. At second order, where we usually stop, the first-order result appears as a driving function, a forcing function, on the right-hand side, driving, again, a linear differential equation. What you have done is created a sequence now of inhomogeneous differential equations that can often be solved by standard linear analysis.

The main point of this slide is to indicate that suppose we had a sinusoid at the source of a single-frequency ω . The first tone you would look for at second order would be a second harmonic component. If you go to higher order, you will see feedback, then, to the primary wave and, in addition, a 3ω component.

If, as one often does, one projects a bi-frequency tone into the system, two tones into the system, you get not only second harmonics but you also pick up sum-and-difference frequencies. It is this latter case that I will show you next in the experiment that is the first one done, really, in nonlinear acoustics for harmonic generation.

(Transparency)

I am out of order here. This is slide 27 in your notes that I have brought forward. I show this to you for several reasons. One is because it was the first experiment. This was by Thurs, Jenkins, and O'Neil at Bell Laboratories in the 1930s. They performed just such a perturbation expansion to solve for these sum-and-difference frequency components.

They then performed their experiment in a long waveguide. The source is here. The wave propagates along. Measurements are taken along the pipe as in Dr. Keolian's experiment. The rest of this U tube is really termination to provide a progressive wavefield.

In each of these cases, what you find is that for the lowest order perturbation, when you solve for these harmonic components, they will grow as a linear function of distance from the source. So these are results for the second harmonics and the sum-and-difference frequencies.

The second reason I cite this experiment is because notice the gap here between theory and experiment. Consistently, for some reason, these authors' results were about 3 dB below the analytic predictions. This discrepancy, in fact, troubled the theorists for a number of years, thinking maybe there was something wrong with the basic equations of motion.

A third reason I mention this slide is to point out that it was Izzy Rudnick --

(The Penn State students stand.)

-- who finally put this problem to rest and the fourth reason was to see if I could get this impulse response from the system, as everyone else has.

(Laughter)

He, in the 1950s -- I do not know if he had gone to UCLA yet or if he did this at Penn State - had a masters student reproduce this experiment, and there is no 3 dB discrepancy. Maybe it was a root 2 or a calibration error; we will never know what happened there. The point is, the curves should sit on top of the theory and in all experiments since then have.

(Transparency)

Suppose we have nonlinearity now, instead of in an ideal gas, considered in an arbitrary liquid. If we are going to be doing analyses based on various perturbation expansions, often we are concerned only with terms of quadratic order. If we have an arbitrary liquid, we start with our generic form of the isentropic equation of state and use a Taylor series expansion in the density:

- the coefficients are then derivatives of the pressure with respect to that density;
- introduce acoustical pressure, excess density;
- rewrite it as an expansion in terms of these small quantities;
- and the quantity that we are after is the ratio of the coefficient of the quadratic term to the linear term.

This ratio, B-to-A, where B corresponds to this derivative, A to this one, is called the parameter of nonlinearity and gives one an indication of the extent to which an arbitrary liquid is nonlinear.

In liquids, this parameter, B/A , varies usually in the range of 5 to 10. For a perfect gas or a diatomic gas one finds that it is $\gamma - 1$, or 0.4.

If we pursue this analysis further, one finds that exactly the same propagation law can be derived, C naught plus βU . In this case, β , our coefficient of nonlinearity, has been generalized to $1 + B/2A$. It is this form of the nonlinearity coefficient that you will find most widely in the literature. If we plug in values, we find that β equals 1.2 for air and 3.5 for water. In this sense, one can say that water is more nonlinear than air.

As a matter of fact, if we remember that this 1 is associated with convection and we see that in air the nonlinearity coefficient is 1.2, one concludes that nonlinear effects in air are due to the kinematic nonlinearities in the equations of motion, whereas in water the nonlinearity is due primarily to the equation of state, so slightly different mechanisms dominate.

(Transparency)

How do you measure this parameter? The formulation on the last slide gives you, by direct expansion, explicit definitions of pressure variation with density for each of these coefficients. However, from an experimental point of view, for a liquid it is a challenging experiment to vary the density and look for changes in the pressure to measure these quantities.

Alternatively, recognizing that the sound speed is the derivative of pressure with respect to density, one can obtain these two results for the coefficients of nonlinearity. In the case of B/A ,

now what we are measuring is variation in sound speed as a function of change in pressure performed isentropically.

The advantage of this formulation of the expression is that one way you can change pressure isentropically is with the sound wave itself. In fact, that is one way that is used, often called the finite amplitude method, of measuring B/A . In particular, one uses one wave to modulate a second wave, and one looks for phase changes in that second wave to deduce this variation in propagation speed.

The second method that is used in the literature is referred to as the thermodynamic method. This comes from reexpanding this derivative from a constant entropy to two derivatives, one performed at constant temperature and one performed at constant pressure. So we have isothermal and isobaric changes now of the liquid and we are measuring changes in sound speed with respect to pressure or temperature.

These experiments are done quasi-statistically and these experiments typically have much more reduced error than those done by the finite amplitude method. In the literature, those quantities associated with the thermodynamic method are, by and large, considered the most reliable. In fact, this comes from Coppen's paper, who was a student of Bob Beyer, now working at the Naval Postgraduate School.

DR. GARRETT: Now retired.

DR. HAMILTON: Was working at the Naval Postgraduate School.

(Transparency)

These tables here are given to you just for reference and to give you an idea of the span of these parameters of nonlinearity, just to give you a bird's eye view that, indeed, for a variety of liquids the coefficient varies primarily between 5 and 10 for liquids, with the exception of, say, for example, sodium and potassium, with slightly less coefficients of nonlinearity.

(Transparency)

The second table I have given you here is for biological media. A large part of the research over the past decade into measurements of the parameter of nonlinearity, B/A , were done by the biomedical community for two reasons.

One of them is because the levels of ultrasound now used for a variety of applications, in particular, lithotripsy, which we will talk about a bit, the disintegration of kidney stones in the body by the passage of shock waves, to model these processes one needs to get a handle on what are these coefficients for tissue.

A second reason for this was that it was proposed that the extent to which tissue is cancerous can be indicated by the amount of water content in the tissue. Healthy tissue typically has a higher coefficient of linearity than water and the water would tend to reduce this factor. There were a

number of proposals that perhaps B/A could be another imaging parameter that could be used for diagnosis within the body.

The idea is that although this is a second-order effect, namely, nonlinear effects are required to measure this coefficient of nonlinearity, they are larger than the effects due to linear propagation of sound through these same tissues, although the latter are first-order effects. The jury is still somewhat out on that.

DR. GARRETT: Are there people doing it?

DR. HAMILTON: Nobody is doing the imaging. It is only laboratory experiments at this point. I am rather skeptical myself.

(Transparency)

Here is Dr. Keolian's experiment of the suppression of sound by sound. This will give you an idea rather dramatically of the way the distortion process takes place to illustrate it. Suppose we have one sinusoid, low frequency, high amplitude, and on top of that at the source we superimpose a high-frequency, low-amplitude sound wave. Now let this wave distort.

Typically, this high-frequency waveform is, for example, several orders of magnitude less in amplitude than the low-frequency one, and you cannot even see it on this waveform. So the effect of the nonlinearity is really to distort the low-frequency so-called pump wave. As this wave steepens up here and expands here, you have a compression of the cycles of the high-frequency wave, where the shock is going to form, and an expansion of the cycles on the other portion.

If you subtract these two waves from each other, what you have, in essence, is a frequency-modulated waveform. In fact, it is phase variations similar to this that are used in the finite amplitude method for measuring these coefficients, B/A , in the so-called finite amplitude method.

For weak distortion, this modulation in phase turns out to be sinusoidal. If you have a sinusoidal phase modulation of a tone, if this is the spectral line corresponding to the high-frequency wave, you get a solution well known in electrical engineering that depends on Bessel functions for all of these sidebands.

The spectral line of interest in the experiment you saw is that of the high-frequency tone at the source. It is proportional to a zero-th order Bessel function and it depends, for example, on distance from the source. As we move farther from the source, the amplitude goes down, ultimately passes through a zero, and then oscillates. This is referred to as the suppression of sound by sound. There is no absorption taking place, energy is conserved; what you have, really, is just a redistribution of energy in the frequency domain.

(Transparency)

Here is an experiment that was done by Jim Tencate for a masters thesis down in Texas. Here is an example. This high-frequency wave is already mixed in with this low-frequency wave (you cannot even see it, so this was extracted by subtraction). As the wave distorts, you see, in

fact, the sinusoidal modulation of the high-frequency wave. As the wave starts shocking up, it just kind of scrunches the cycles at the point of shock formation. That is exactly what you saw.

I will skip the spectrum; you saw that in the experiment.

(Transparency)

I will show you the world's record for suppression here. This was done by a group in Nanjing, China, in fact, where the International Symposium on Nonlinear Acoustics was held just last week. They achieved some 47 or 48 dB of suppression of the high-frequency tone.

(Transparency)

If one takes a look at these waveforms from that analytic expression that I showed you that was derived from Poisson one can calculate the derivative of the particle velocity with respect to time as a function of point on the waveform. The question you want to ask is at what point in space does this wave form a shock?

The point at which we say shock formation has occurred is when a vertical tangent has appeared in the waveform. This calculation is fairly simple to do. What it does is yield the following result.

(Transparency)

We call this distance \bar{x} and it is very simple, it is $1/\beta \epsilon k$ -- β the coefficient of nonlinearity. ϵ , if we are talking about sound waves in fluids, is the acoustic Mach number at the source. If we are talking about compression waves in solids and progressive plane waves, this is equivalent to the strain, so this would be the strain at the source. k is the wavenumber of the frequency of the source. This is assuming a single-frequency source with the waveform going into shock.

It is not surprising that this shock formation distance will decrease as the nonlinearity goes up or as the source amplitude goes up. It may not be quite so obvious why it is inversely proportional to frequency. That happens because for the same amplitude wave, as I increase the frequency, I have increased the slopes of that wave at the source, so these waveforms need not go so far to distort to the point that shock formation occurs. So shock formation becomes more likely at higher frequencies.

DR. GARRETT: Another way to look at that is that the natural length scale in the problem is the wavelength, so they still go the same number of wavelengths to form a shock.

DR. HAMILTON: That is right. Thanks.

The way I have got this written in this slide is to show you that all of these principles can be extended to the case of, in this case, an isentropic solid. Whereas we had a coefficient, B , to describe the quadratic nonlinearity in the equation of state for an arbitrary liquid, for an isentropic solid there are three such constants that are necessary.

The notation I have used here, A, B, and C (you will see how they appear on the next slide), are the so-called Landau coefficients. Various authors use various notations. There is another popular one that you may be familiar with, the Myrnam notation. The point is you have three constants describing the state at that order.

The main point here is to make calculations of the coefficient of nonlinearity for a variety of materials. If we look for air as being a representative gas and water being a representative liquid, and iron, just to pull one for a representative isotropic solid, we see that the coefficients of nonlinearity for each of these materials is of order unity.

However, once you start introducing microfeatures into the materials, for example, voids due to foams or microcracks inside rocks; if you have a very stiff material, say, a rock, and you have cracks inside and you compress this material, some cracks open, some cracks close, and the dependence of the stiffness on the amplitude of the compression is highly nonlinear. This nonlinearity coefficient can become very large due to inhomogeneities, grain-grain contacts, crack formation, delaminations, for example, in materials, and beta can take off by orders of magnitude.

For this reason there is considerable interest right now in using or developing methods for measuring third-order elastic constants, or nonlinear effects, in solids, because information from these coefficients can tell one something about fatigue states and stresses within these rocks.

For example, Schlumberger is very interested, because they have determined that large coefficients here can indicate that rocks are close to fracture for the point of oil well exploration.

(Transparency)

This is to show you the equation of state for an isotropic solid. This is the strain energy density -- μ is shear, K is bulk modulus. This is where your Landau coefficients appear.

DR. KEOLIAN: Is anyone using this to predict earthquakes?

DR. HAMILTON: Not to my knowledge.

It is interesting -- these are called Landau coefficients -- where this equation appears is in a homework problem in Landau's volume on elasticity.

(Transparency)

Let us look at another type of wave. How do shear waves distort? There is no compression now. If you work out the equations for perfect shear in one-dimensional progressive wave motion, one finds that the nonlinearity is no longer quadratic, but it is cubic.

One can then develop a propagation law that is not dissimilar to that for liquids and gases for sound waves, except that instead of beta times U particle velocity, it is beta times particle velocity squared. You have taken away the sine dependence now of the motion of the particles in the medium; that is, the compressive parts still steepen forward, but now so do the negative ones. So the prediction of a model like this would be an asymmetric distortion following, however, a similarly simple propagation law.

In this case, with the nonlinearity cubic, as that wave propagates, instead of generating all harmonics, all integer multiples of the fundamental source frequency component, now all you pick up are the odd harmonics in this distortion spectrum.

(Transparency)

The final waves and solids that I want to just give passing notice to are surface waves. Surface waves in the context here are disturbances that propagate along an interface. Rayleigh was the first to derive an expression for the linear propagation of surface waves when you have an isotropic half space bounded by a vacuum.

Various extensions of this theory considered different loadings. For example, Stoneley considered the case where you have two isotropic solids in contact and you have an interface wave propagating along. The particle motion associated with these surface waves is far more complicated than that in just a compressive or a shear plane wave.

In the lower medium, let us say the host medium, even for a Rayleigh wave, near the surface you find that the particle motion is much like in a water wave, except that it is retrograde. If the wave is going from left to right, the particle motion rolls in the opposite direction. As you go deeper within the solid, the direction of that rotation changes direction, changes polarity, but the disturbance is confined to a depth on the order of one wavelength of that disturbance. And this occurs for a Stoneley wave on either side of that interface.

A Scholte wave is what happens if you take the Stoneley wave, essentially, and set the shear modulus equal to zero and turn the upper medium into a fluid. The Scholte waves are of particular interest to the oil exploration people. These are called tube waves in the low-frequency limit. If you think of a bore hole, this is filled with liquid, and you have a solid outside.

At high frequencies, if you would have a liquid inside and a solid outside, you will get a Stoneley wave, as you see here. As you reduce the frequency, the wavelengths get longer and inside it looks more and more like a plane wave and the roll-off into the surrounding solid becomes more and more gradual. These are called tube waves.

(Transparency)

What does the nonlinearity in these waves look like? The distortion process is something like this. The distortion is substantially different from that in compressive waves. The reason has to do with the properties, actually, of linear surface waves. For example, if you imagine this Rayleigh wave propagating along the interface at a frequency ω , it distorts and generates a second harmonic component.

That second harmonic component has a wavelength that is half that of the fundamental, so its depth penetration moves upwards. Higher harmonics are generated that are trapped closer and closer to the interface. What you find is that the prediction is that as the propagation of this wave moves forward, the energy density actually increases somewhat at the interface and you get a rather

unusual distortion sequence at that point. Instead of a simple sawtooth waveform, you get one with cusping.

You have to be a little careful in surface waves. There are two components of the particle velocity. You can talk about an in-plane component, horizontal, and a vertical component, perpendicular to the interface, and these waveforms are considerably different.

(Transparency)

We have taken now our simple propagation law and marched our waves up to the point of shock formation, but we have not described at this point how to model the waveforms where these shocks occur. That is the second part that we will begin.

(Transparency)

In fact, this was the main difficulty following the analyses in the 19th century by Stokes and Poisson. The thermodynamics, the understanding was not there until the end of the century, when Rankine and Hugoniot developed relations for shock waves, step shocks, in fluids.

In particular, what they did was they assumed a perfect jump discontinuity in the fluid. They said, "Okay, we know the state of the fluid behind the shock -- pressure, temperature, internal energy, density, et cetera -- and we know the state in front of the shock. If we now draw a control volume around the shock, then we can do momentum, mass, and energy balance across the shock."

Rankine and Hugoniot thus derived relations which give us propagation speeds and various other parameters related to the propagation of that shock. It was very difficult to incorporate these equations into the nonlinear theory that I showed you initially. In fact, this stymied people. Despite the fact these equations were derived at the end of the 19th century, it was not until about the 1950s that people understood that to make certain approximations you did not need the full artillery of the Rankine-Hugoniot shock relations to describe the propagation of sound.

Much lower amplitudes permitted approximations that allowed one to marry these two theories, and this is called weak shock theory. It was developed by people like Landau and Wideman and others during the 1950s.

What we assume is that we have a waveform that at some point became sufficiently steep as to form effectively a discontinuity in that waveform. There are two parameters now to describe the propagation of the shock: one is its position and the second one is its velocity.

Let us consider the velocity, first. If you take these full shock-wave solutions and do expansions and consider only the lowest order terms, you find a very simple rule that tells you how fast that shock wave is going to propagate: The shock wave will propagate not at the small-signal sound speed, C_0 , but at that speed corrected by the following quantity. Beta is our nonlinearity coefficient, again, and what it multiplies now is the average of the particle velocities just behind the shock and just ahead of the shock.

This result is consistent with our nonlinear propagation law for the continuous distortion of the waveform, because as that shock amplitude vanishes, point B approaches point A, and this average becomes simply the same value, that of the particle velocity in the fluid at that point.

The way this simple propagation law is married with the previous analytic results is we use the Poisson solution to describe the value of the wave behind the shock and before the shock, and you can think of weak shock theory, then, as a matching condition, or a boundary condition, to marry these two solutions. Frequently one can obtain, then, analytic solutions based on this simple law.

Landau was able, on the basis of this weak shock result, to derive a very clever geometric method for determining the location of the shock. This is referred to as the equal area rule.

(Transparency)

Essentially by doing an integral across this wave, he said, "Well, let's just let the Poisson solution take this wave and become multi-valued, although physically realizable, and knowing this part of the wave is physically appropriate, and this part of the wave is physically appropriate, at what point should we put the shock to give us a physically meaningful result?"

It turns out that application of that weak shock theory tells you that you just take the Poisson solution, permit it go multi-valued, and place that shock at precisely that point where this area that it cuts off is equal to this area.

First of all, it gives you a lot of intuition as to where the shock would be when you are looking at various solutions and, second of all, one can derive analytic results this way quite readily for some simple cases.

The elegance of weak shock theory is that whereas I told you in the beginning of the lecture it is the viscosity and the heat conduction that is maintaining our shock front and offsetting the tendency of the wave to steepen forward, you do not see any viscosity or heat conduction terms in these results. So you are able to essentially describe the shock without having to know what your viscosity is, what your heat conduction is, what your bulk viscosity is, for that matter.

The price you pay is that you have assumed that this pressure jump is a perfect step discontinuity, so this result gives you no information whatsoever about shock thickness and, therefore, rise times associated with shocks, which can be of interest, for example, in sonic boom analysis. Nevertheless, it is a very convenient mathematical tool.

(Transparency)

This is a page out of Rudenko's and Soluyan's textbook on nonlinear acoustics, a very widely used book that was published in the 1970s. I give you this just because these two paragraphs walk you through the geometric argument for taking a triangular wave, steepening it forward to kind of the beginning of a sawtooth, letting it continue to steepen, chopping off the equal areas, and what they do, just by geometric construction now, is they work their way down

and say, "Okay, we can calculate length AE" -- AE is here, if you will, the duration of your pulse -
 - "and calculate line segment ED," and that gives you the amplitude of this pulse.

The solution that they give you is effectively the solution for the N-waves that we saw in the experiments on Sunday. What we find here is that the amplitude of the N- wave, ED, goes as roughly 1 over the square root of the distance of propagation. So the amplitude of this N-wave decreases with distance because of the energy being dumped and being dissipated at that shock.

(Transparency)

Some examples of N-waves, just to give you some ideas of orders of magnitude of time scales involved:

- . For sonic boom work, durations of order 1/10 of a second, rise times of order milliseconds;
- . spark sources in air -- David Blackstock has done a lot of these experiments in Texas for modeling purposes -- here we have rise times on the order of microseconds;
- . and lithotripsy pulses, those that are used to disintegrate kidney stones, have rise times of order nanoseconds.

The point here is that in each of these cases the rise times are orders of magnitude smaller than the durations of the waveform and in these cases it is appropriate to use weak shock theory to describe the propagation of these N- waves.

I might note in passing that in the experiments that have been done in Texas by David Blackstock and his students, in measuring these spark sources they construct their own condenser microphones. You do not go out and buy megahertz microphones for measurements in air; you build them in your laboratory.

Then the question is, how do you calibrate them? Where are your calibrated sound sources? He uses the N-waves themselves. The propagation of this N-wave is a nonlinear function of distance and it depends on, simply, material properties and amplitude at the source; it does not scale linearly.

If you double the input, you do not double the output. By matching the output curves with the input voltage, he has an absolute calibration method for determining what the amplitude of the sources is. You can actually use nonlinear propagation laws to give you an absolute reference point, provided you know everything about the fluid: nonlinearity coefficient, density, temperature, and so forth. You can do these calibrations absolutely.

(Transparency)

In case you are unfamiliar with the device, here is a picture of a shock-wave lithotripter. The lower half of the picture is the version more commonly in use today. There are several ways of generating the shock. The one made by Dornier Company in Germany is to use -- basically, they put a spark plug in water and fire the thing off and focus the thing with a mirror.

This mirror is an ellipsoid, a metal ellipse, and if you can imagine constructing an ellipse that includes the location of the spark source as one focus and your kidney stone as the second focus, then the waves generated at this point will reflect off here and converge toward the second focus.

You are coupled by a water bath in between and they learned quite quickly that you needed to put a water bag on the exit side of the body, because the tissue-air interface makes a very good pressure-release surface. So when those waves would hit the tissue-air interface, they would phase-invert and now you had a many megapascal negative-pressure pulse acting on your blood vessels, and they were cavitating and bruising due to the reflections internally from the body. What they do is they postpone the reflection by putting water bags here and the reflection now takes place up here.

These lithotripters are actually used now, I have learned, on an outpatient basis. They are actually blood mobile-type vans that can roll up into your driveway, park you inside, blast you for two or three hours, and put you back on your couch to recuperate.

DR. GARRETT: Do they blast you for that long or does it just take that long to get off the table?

(Laughter)

DR. HAMILTON: All I know is they are doing it on an outpatient basis. You do not spend the night in the hospital. And they have done it in homes. The only reason I know about this van is I saw it on the news in Texas; they were talking about one up in Dallas that was --

DR. GARRETT: But it is not a long procedure, right? They just blow up your kidney stones and then --

DR. HAMILTON: They let you pass them, that is right, so then they are in your urinary tract and presumably the particles are small enough that, without such extraordinary pain, they pass through.

DR. GARRETT: It takes thousands of shots.

DR. HAMILTON: It takes thousands of shots. My understanding is it takes a couple of hours. And there are considerations like timing the shots with your heart rate, you do not want to get into any body resonances and things like that, so there are ancillary concerns involved with this.

This is one type of source. They also use piezo- sources that can, I guess, give you tighter focuses than these spark plugs will. I guess there is another company that uses a kind of capacitive discharge. They charge a plate and just let the thing go and hit you with a shock that way.

(Transparency)

Here is a picture of one of these pulses. Here is a waveform fairly near the source, and the distorted waveform here. Again, these overpressures can be on the order of 50 to 100 MPa. Incidentally, the mechanism that seems to be apparently or primarily responsible for the stone

breakup is cavitation. The bubble collapses that Dr. Atchley showed us, as these collapse they emit shock waves -- if they can chew up propellers, they can chew up kidney stones. This appears to be the mechanism that is generated by these high-intensity shocks.

Continuing on to the next part of the lecture, there are some explicit solutions that can be derived from the equations that we presented.

(Transparency)

In particular, there are some very widely used, explicit frequency-domain solutions. If you assume a source condition that is sinusoidal -- τ appears again and again, but I have not defined it -- τ is a retarded time. So we put ourselves in a time frame that is moving, really, with the zero crossings of the waveform.

One can invert the implicit Poisson solution and derive an explicit spectral solution, and it appears in a Fourier series like this. The spectral amplitudes now are functions of distance from the source and we will use a dimensionless distance parameter, σ , that is X over the shock formation distance, for that sinusoidal disturbance. $X = 1$. We have just formed a shock at the zero crossings in the waveform.

In the pre-shock region, σ less than 1, a solution was derived by Fubini in the 1930s in terms of these Bessel functions. In fact, Logan Hargrove rediscovered this in the 1950s. A second solution occurs as a limiting case farther from the source, when we are beyond about three times the shock formation distance. The waveform is substantially sawtooth in appearance and the limiting form of the coefficients is this.

DR. CHAMBERS: Is the retarded line X/C_0 or X/C_0 plus βU ?

DR. HAMILTON: No, X/C_0 . It is the linear retarded time. You can think of yourself as moving with zero crossings, any zero points on the waveform. There is an integral solution that couples the regions 1 to 3 that I have not shown here.

This is what it looks like, so if you launch a tone at the source, it has a single spectral line, and the amplitude of the fundamental component decreases because it is pumping energy upwards in the spectrum and feeding these higher harmonic components, namely, the second harmonic, third harmonic, and all the others start growing.

These grow up until roughly the point of shock formation and now nonlinear losses start taking their toll and stripping away at the shock. Eventually, when we are several times the shock formation distance, we are in this sawtooth region. The harmonic amplitudes are related as $1/N$ for any sawtooth wave, and they decrease with distance roughly as $1/X$.

So at distances far from the shock formation distance, the amplitude of the sawtooth wave is going off as 1 over distance, as though it were a spherical wave propagating in free space. However, this is a plane wave and the energy loss is irreversible; it is being dumped into the shocks.

STUDENT: Excuse me, is that just the general function?

DR. HAMILTON: It was the particle velocity.

For any of these results I am showing you, now that I think of it, you can use the plane wave-impedance relation to reexpress this in terms of the pressure with no loss of accuracy in the derivation that led you to this point. This could equally well have been a solution for the acoustic pressure as a function of distance in, say, a waveguide.

(Transparency)

An interesting thing happens if we examine the sawtooth limit of this solution a little more carefully. We have rewritten it here with the Fourier coefficients now expressed explicitly for the sawtooth. Remember what σ is: It was distance over shock formation distance. This distance parameter, therefore, is proportional to the particle velocity at the source, U_0 . So U_0 is your source amplitude.

Now, suppose we are far beyond the point of shock formation? In this case, σ is much greater than 1. We can throw the unit away and replace with the dimensional quantities, and look what happens. The source amplitude, U_0 , drops out of the equation. So the solution that we get here is completely independent of the amplitude of the wave that was launched into the medium. This is true for all of the harmonic components.

That means, if you think of ways of varying σ , because this is a large σ approximation, there are several ways to do it. The two most obvious are either moving away from the source as a function of distance or, more commonly, from the point of view of, say, an experiment, you have got a microphone at a fixed position, X , and now you are going to increase the source amplitude and watch its effect take place.

Let us look at the fundamental component for a second. This happens, as you see, for all of the spectral components. Let us look at $N = 1$. You can imagine an experiment as I have described. We have the source here. We have the receiver down here. Now we are going to turn up the amplitude.

If the amplitude is low, we are in the linear region of the system. If I put in 6 dB I am going to see an extra 6 dB come out the other end. As I increase it higher and I put in, for example, double the amplitude, again, and another 6 dB, I may be seeing only a 3 dB increase coming out the other end. Why? Because energy is being lost at the shocks and I am looking at the fundamental component.

Eventually we hit what is called acoustic saturation. No matter how much more energy I am pumping in from the source, I am not going to see any more at this observation point. How can that happen? What is happening is, as I continuously increase the source amplitude, and I have got a fixed distance, L , between source and receiver, the point in between where shock formation is occurring is moving closer and closer to the source.

Therefore, that wave has a longer and longer propagation path in which to dump its energy -- basically you are heating the water -- between the source and the receiver. You are just moving shock formation closer to you, things cancel and balance in just such a way that you see saturation on this end.

This happens for the case of a sawtooth wave. It will not happen for an N-wave. N-waves do not saturate the same way. As I mentioned, a sawtooth wave decays as 1 over distance because of energy loss at the shock. An N-wave decays as 1 over square root of distance due to energy loss at the shock. It is not fast enough to give you this full saturation effect.

DR. SWIFT: Could you make any comments or speculations about the spherically imploding shocks as we have in sonoluminescence and whether the shock heights grow as the spherically imploding shock wave implodes? Effects like this, do they occur?

DR. HAMILTON: Before all the interesting phenomena are happening, you are getting the same losses occurring at the shock fronts as you are converging, but as soon as things get interesting, this entire analysis is going to fail.

DR. SWIFT: But this is a plane wave analysis and right from the start you have got a spherical

DR. HAMILTON: In fact, I was just talking with Hank Bass about this a couple of days ago. Although, for example, although you have these Rankine-Hugoniot relations for an arbitrarily large step shock in a fluid, for plane waves one does not have an analysis like that, even for spherical waves. It is not so simple.

For example, for underwater explosions, most of the data you have on either diverging or imploding shock waves are, for the most part, empirical. I have not seen, certainly, anything analytical with respect to that.

(Transparency)

Let us take a look at an experiment. This is one done by David Blackstock and his student, Don Webster, down at the University of Texas. In this case they have a 30-meter waveguide. They input, in this case, a signal of 500 Hz and looked at the second harmonic, looked at the third harmonic, the point being -- no surprises here -- all of these harmonic components saturate as a function of source sound-pressure level. Here they are raising the sound-pressure level in air up to about 160 dB.

The other curves on here are variations of other parameters, distance from the source, frequency of the excitation, and so forth. But the main effect is here.

(Transparency)

The other experiment I will show you was done by Jack Shooter and Tom Muir with David Blackstock down at Lake Travis in Austin, where a lot of the sonar experiments are done. David Blackstock tells the story, he was teaching nonlinear acoustics, I guess in the early 1970s, and he

had Tom Muir in the class and gave the lecture on saturation. Tom said, "I don't believe it." So he hauled his hardware out to the lake and dropped his sonar transducers in there.

These are the data he got in a now well-referenced paper. These are beam patterns due to radiation from a source 3 inches in diameter radiating at 450 kHz, roughly half a megahertz. These beam patterns are taken at increasing distances from the source.

The lower row is for low-intensity excitation of the source. The upper row of beam patterns is for high-intensity excitation of the source. In the lower row, what you see are the formations at about four yards -- I guess the near-field collimation length for this beam is on the order of a yard or two, as I recall. The side lobes start becoming well-defined at this point and they are about, according to linear theory, 17 dB or so below the main lobe. The beam pattern stays this way with distance.

In the upper row the picture is substantially different. As the wave propagates, we see that apparently the side lobes are increasing in this beam pattern. What is happening here is that the nonlinearity is acting most strongly in the center of the beam where the wave amplitude is highest. The center of the beam is going into saturation.

What is happening is, if you think of the beam pattern, the amplitude is the highest in the center, that is going to saturate first, then the edges, so the main lobe in that beam pattern is going to become more and more blunt. The side lobes, which are already starting out with a deficit of 17 dB, have virtually no nonlinear effects taking place there, so they are well-described by linear theory.

Instead of the side lobes coming up, really, it is the main lobe coming down. It is flattening out and its beam length increases. This blunting is a well-known effect in diffraction patterns at high intensities.

(Transparency)

Suppose our weak shock theory is not sufficient to describe all the processes that might be of interest. For example, suppose we want to take into account various dissipation mechanisms. Blackstock and Pestorius came up with a numerical algorithm for doing so, which is a kind of hybrid scheme that takes advantage of the simple laws of propagation that we introduced in the beginning of the lecture, namely, a propagation speed law for continuous segments of the wave and a propagation speed law for the shock waves joining these continuous segments.

What he does is, he says, "Okay, I've got a waveform here at the source. I'm going to propagate it forward ΔX . I'm going to steepen the wave and if I see a shock, I'm going to correct it according to those propagation laws." Now he says, "Well, I may have attenuation and dispersion that need to be taken into account, so let's go back and repeat this step and, this time, put in these effects."

He includes these effects in the frequency domain. So he comes back, takes a Fourier transform, attenuates the wave, and now that he has got all the spectral lines separated, he can apply the dispersion and phase shift them, reassemble the wave at the end of delta X and goes on to the next step and marches forward.

The advantage of this procedure is it takes advantage of the modeling of the shocks in the time domain and in the frequency domain, it permits you to put in any kind of an arbitrary attenuation or dispersion law that may happen to suit the system.

This particular algorithm, in fact, is still being used. It is being used now at the University of Mississippi. It is one of the ingredients of a computer code for modeling sonic boom propagation through the atmosphere. In fact, there is a paper about to come out, and most of the authors are here. Hank Bass, Rich Raspet, and Jim Chambers are using a code based on this to model sonic boom propagation.

(Transparency)

Let us take a quick look at how it works: first, kind of a test run; computation of a sinusoid in the waveguide. The source level is 160 dB, the frequency of the excitation is 800 Hz, and now we watch the propagation of this sinusoid down the pipe.

Weak shock theory would tell us that the waveform should have a sawtooth shape, given by the dotted line here. What we find is that the proper solution must take into account boundary layer effects that we heard about in Professor Bass' lecture. The shear at the wall and the heat conduction at the wall give us both losses and dispersion. These are substantial at these frequencies in waveguides.

The absorption reduces the amplitude of the wave by about a half and the dispersion gives you the rounding and the slight phase shift and asymmetry in the wave. These are old oscillograms from the 1970s in Mike Pestorius' dissertation, but the theory and experiment sit right on top of each other.

It was not for modeling simple sinusoidal propagation that this code was developed. It was developed to analyze noise signals.

(Transparency)

What happens when we have a random waveform? The basic laws really do not change. Suppose we start out with a jagged straight-line segment, given by this figure. Now let us apply our propagation laws. The Poisson solution tells us that the excess propagation speed is proportional to the local wave amplitude. So I double the amplitude, I double the amount that is going to advance forward.

The upshot is straight-line segments will remain straight-line segments. So this wave just starts pivoting like this. We form shocks, and now these shocks are going to propagate at different speeds. Weak shock theory tells us that because the excess propagation speed is related to the

mean particle velocity, this higher shock is going to propagate faster. This shock is going to propagate faster than this shock. Therefore, this one will catch up.

What is expected to happen is, as they catch up, they finally merge. So shocks take over shocks. You have an irreversible alteration of the wavefront and that information is lost. You can anticipate what is going to happen here.

If you start off with a stochastic signal at the source, a random signal, and let the wave propagate, first the wave is going to start forming shocks, so energy is going to be pumped upwards in the frequency spectrum. Now you have a distribution of random shocks in the waveform. These shocks will start catching up with each other and, therefore, you are reducing the number of shocks in the waveform, the spatial frequency of the wave decreases, and energy pumps back down in the spectrum.

First it shifts upwards during shock formation and then shocks start eating other shocks and you get less of them per unit of time and the energy cascades back down, again. Let us take a look at that.

(Transparency)

Here are the experiments corresponding to that model, waveform measured near the source and used simultaneously for the input to the computer code, and let us follow the measured waveform as it propagates. As we go from here, a distance of 1 foot, to 13 feet, we see the steepening of the waveforms and the shock formation.

Let us follow, for example, this bump here in the center. As we propagate, this shock now is going to catch up with this lower amplitude disturbance. Here it is, again. By here it is gone. This is the process of these shocks running over each other in this alteration of the spatial frequency of the wave.

(Transparency)

They took frequency spectra corresponding to this experiment to show that this is happening. Here is the frequency spectrum at the source. As the wave propagates to 50 feet we have the dashed line. Frequencies are extended to both higher frequencies and to lower frequencies, so energy is going both upward and downward on the spectrum simultaneously.

Why is it going down? You can think of shock formation, again, as increasing the frequency components of the wave. Another way is to think about this whole spectrum as being a sequence of spectral lines that are all multiplying quadratically times each other.

All of the spectral lines in here multiplied together would give us not only some frequencies on this end but their differences will contribute to a low-frequency wing of the spectrum. Now, as the wave propagates, we go out to 73 feet. Linear attenuation mechanisms start dominating the nonlinear acoustical mechanisms. The wave spectrum becomes increasingly low-pass filtered and

energy is stripped away from the high-frequency spectrum, but energy is still to some extent pumped into lower frequencies, where the attenuation is much lower.

(Transparency)

Part of the motivation for this work was in the interest of aircraft noise. In aircraft noise often you have very strong tonal components mixed with a background broadband noise spectrum. Here is an experiment to demonstrate that nonlinear mixing.

What we do in these three rows of frequency spectra is consider, first, a high-intensity tone propagating alone. Now turn off the tone and consider low-intensity noise propagating and then we will mix them together. So we start off with the tone at the source at about 150 dB. Very close to the source we still have harmonics (some of this may be nonlinearity due to the driver itself, some of it perhaps due to the short propagation path), but as the wave propagates quite quickly, it is distorting and the amplitudes of these harmonics grow up in correspondence with the equations we have been looking at. So we are well beyond shock formation at this point.

Turn off the tone and look at the noise. We now put in broadband noise between about 1 and 3 kHz, with a roll-off here, and let this propagate down the pipe. Not a lot is happening here, certainly no nonlinear effects. The only main effect is attenuation of the wave at the high-frequency ends, so the higher frequencies are being attenuated slightly more than the lower frequency components.

Now let us mix the two together, tone and noise, at the source, and here we see, at the source, the mixing due to the noises as we saw without the tone. Very quickly, though, during the propagation process, this energy in the noise band is pushed all the way through the high-frequency spectrum that has been generated due to the harmonics in the finite amplitude tone.

How is this happening? What we have got now is a very broadband redistribution of the noise in this low-amplitude waveform, all of this evidently happening due to the finite amplitude effect of this high-intensity tone from the source.

(Transparency)

This can be seen by a numerical model using the computer code that I just described. Let us start with the tone and a simple distribution of noise here between 1 and 3 kHz. Now think about what is happening in terms of spectral interactions. This tone is going to interact with all of the spectral lines in this pedestal here.

It will push frequencies out in the first approximation all the way up to the sum frequencies of the highest frequencies in this band. If I have this frequency and this frequency, I end up out at a point roughly here. The interaction of this tone with the band to the first approximation pushes out the first pedestal here.

This tone can now interact with this pedestal and generate some frequencies out here. So this high-intensity tone is continuously multiplying and mixing with the spectral lines that it itself is

generating by the interaction of itself with the noise band and progressively pushing the noise outward, all the way out to the full extent of its own spectrum.

(Transparency)

The purpose of the next few slides is to just give you what are the most widely used model equations in nonlinear acoustics for describing the propagation of sound, really without derivation, just for reference and just to make a few passing comments about to show you where they come from.

What happens when you get rid of these approximations to plane waves and what happens when we put the viscosity in and the heat conduction, how do you do the full problem?

(Transparency)

The equations describing the system, one can take these, for example, the notation I am using is right out of Landau and Lifshitz: conservation of mass, momentum -- here we have our bulk viscosity coefficient, which was discussed the other day, shear viscosity coefficient -- in the entropy equation, S is the entropy, and here κ is the coefficient of heat conduction and, finally, we have an arbitrary equation of state, which I am writing as pressure as a function of density and entropy.

In the case of a perfect gas you can have an explicit relation for the dependence of the pressure on entropy and density, but that is not necessary, we will not really be using that.

The reason I have highlighted some of these terms in red is that the ordering scheme we are going to use is to categorically dismiss those terms as third order, which I will explain in a moment. We are just going to focus on what we are going to call second-order terms in these equations.

If I look at this term, I say, okay, I pick up two orders because I have got a quadratic field variable, so it is nonlinear, times a small loss coefficient, so I call that third order and throw it out. Likewise, for this shear viscosity term in the entropy equation this is quadratic times a third quantity, which gives you third order. That goes out.

The other term that I will be throwing out to get to our final equations is this one involving the vorticity. The curl of the particle velocity gives me the vorticity. If I assume a fluid with no boundaries, which I will for the purpose of these equations, and if I assume no vorticity, to begin with, then no vorticity will be generated in the absence of interaction with boundaries, so I can take that term out. Really, the equations I will be dealing with will be these.

A question was raised earlier by Dr. Atchley about what happens if you do not have progressive waves and is there an exact wave equation like looking result (sic) for the fully nonlinear problem? There is only one equation I know of that is exact in terms of a single field variable for the case of a perfect isentropic gas.

(Transparency)

That is the equation that I give on this slide. What you have to do is assume no losses in the fluid, assume a perfect gas, and you can write something that is beginning to look like a wave equation in terms of the velocity potential ϕ [phonetic], scale of velocity potential ϕ .

Now we recognize our compound linear wave operator on the left-hand side, but what we see on the right-hand side are both quadratic and cubic terms that must be dealt with.

The advantage of this equation is it is exact. The disadvantage is, except for those very reduced progressive plane wave analyses I showed you earlier, there are no known exact solutions of this equation. At some point, what you have got to do is come back to this point, make some types of approximations, do perturbation expansions, or something, and at that point you have destroyed the exact nature of this equation and you might as well start with something more simple, to begin with.

It is of limited use, in that sense. On the other hand, it is a common starting point, especially in air elasticity. In papers you will see, on perturbation techniques in nonlinear acoustics, oftentimes that is their starting point so they can show exactly where the approximations come in.

(Transparency)

The ordering scheme that gets you tremendous mileage in nonlinear acoustics in terms of asking where do I start with all of these equations, and one of the rules governing the types of substitutions I can make, this point of view began with Lighthill in a classic paper he wrote in the 1950s on solutions of the nonlinear equations. He basically introduced the ordering scheme I just outlined on the earlier viewgraph.

You say that all lossless linear terms are characterized by the acoustic Mach number, particle velocity over sound speed. We introduce another small parameter to now characterize the loss coefficients in the fluid, bulk shear velocity and heat conduction.

Now we start ranking terms. Each of these we assign an order one to. Anything of order epsilon, namely, a linear in the acoustical field variable with no loss coefficients, we call this first order. There are two varieties of second-order terms that come in. One is quadratic nonlinear terms with no loss coefficients, or linear terms with loss coefficients. We now categorically dismiss all third- and higher order coefficients.

DR. RASPET: But that μ squared term --

DR. HAMILTON: You mean μ squared without an epsilon?

DR. RASPET: Right.

DR. HAMILTON: But then there is no field variable. Epsilon is proportional to whether it be temperature fluctuation or density or sound wave. So right, you will not have a μ squared term here.

The corollary to this rule is as follows. If I am going to discard all third-order terms, in my second-order terms I can make first-order approximations. I can take a linear relation, substitute it

into a nonlinear term, because any additional accuracy in that substitution would have been third order and we are dropping those anyway, so this gives you a very nice, well-defined ordering scheme for combining and eliminating variables in these equations.

(Transparency)

To show you what these equations look like in that reduced form, I have presented this slide, where here is our mass equation, momentum equation, we have combined equations of entropy and state. In each of these equations, all first-order terms are on the left and all second-order terms are on the right, and we have thrown out the third orders and we have made a bunch of substitutions to simplify.

The only point I want to remark upon here is this new factor we have introduced, called the Lagrangian density. This is kind of interesting from the following point of view. Suppose we consider a progressive plane wave. At first order, pressure is proportional to the impedance times the particle velocity -- P equals ρCU . That is an order one expression.

I can take this order one expression and it substitute it into the Lagrangian density, because it will still be correct to second order. When I put the plane wave impedance relation in the Lagrangian density, these terms cancel. Therefore, these red terms in both of these equations disappear. Therefore, the momentum equation is linearized.

What this says is to dominant order in the nonlinearity for progressive plane waves there is no contribution from the nonlinearity in the momentum equation and we might as well have started out with a linearized momentum equation.

Where the nonlinearity comes into play for our progressive plane waves is really at two places: first, the conservation of mass equation and, second, the equation of state. In fact, when you look at these two terms, here is your $1 + B/2A$, and there is your β . So there is your coefficient that comes into the progressive plane wave results.

(Transparency)

The equation I want to focus on here is the Westervelt equation. This is perhaps one of the most widely known equations that you will see in the literature. The one above it is the complete equation with all second-order terms included. Δ is our loss coefficient. This includes the shear viscosity, the bulk viscosity, and the heat conduction.

Notice that they combine additively in this approximation, so you cannot distinguish these effects. These are what give rise to the classical thermoviscous attenuation coefficient that we saw in the earlier lecture.

The Westervelt equation here has a single term for the loss coefficient and a single term for the nonlinearity. This equation is widely used for the analysis of, for example, nonlinear sonar systems.

What we want to do, though, is get a more simplified equation for investigating the propagation of planar type waves where we now include viscosity. We do not want to have to appeal to weak shock theory. What we can do now is approximate this Westervelt equation as follows (again, I just want to hit the highlights here).

(Transparency)

The idea is this: If we take a look at the Westervelt equation in the top row and take two limiting cases, first throw out the nonlinearity, set beta equal to zero. A good approximate solution for plane waves is simply the exponential attenuation of a sine wave, where this group of parameters here, as we recognize, is our attenuation coefficient.

The second equation here is our Poisson result I showed you earlier. The only key point I want to show here is the parameters come in the following pairs: We have a retarded time, as I showed you here, and we have the propagation distance, Z , always multiplied by one of our small parameters. In this case, it is Z times the acoustic pressure for the nonlinear solution. Here it is Z times the loss coefficient for the linear attenuating solution.

That is to say, in both limiting cases the field variables have the general dependence of dependence on time in this retarded time combination and, secondly, on space, where distance is multiplied by a small parameter. If we take these as our explicit variables, do a variable transformation, throw out all these third-order terms, again, we find out that this equation at the top reduces at the bottom to something far simpler, referred to as the Burgers equation.

The Burgers equation is a progressive wave equation that describes the propagation of waves, in this case, in the positive x direction, and accounts for thermoviscous losses through this term, and it accounts for quadratic nonlinearity in this term. It is the simplest model equation that accounts consistently for these combined effects and is probably the most widely studied equation in nonlinear acoustics.

Burgers derived this equation, I do not remember when, in the 1940s, perhaps, to investigate turbulence, primarily, and it was rediscovered later as an appropriate model equation, also, for the propagation of finite amplitude sound in viscous fluids.

A remarkable property of this equation, which we will not go into here, is that although it is nonlinear, you can linearize it. You can actually introduce a nonlinear variable transformation, turn this into a parabolic heat diffusion equation, use your linear solutions to solve it, and then invert back.

The solutions are enormous, often unwieldy, but they are exact, and a lot of analyses are based on that transformation.

(Transparency)

Once we are in this form, what we often do is solve this equation numerically. What I have given you here is a very simple algorithm to accomplish that task. Frequently this equation is

solved using a spectral approach in the frequency domain. We are going to non-dimensionalize this equation, divide the pressure by a characteristic pressure, introduce the same dimensionless distance that we had before, distance over shock formation distance.

Here is our retarded time, now scaled by frequency to make it dimensionless and, finally, we introduce this parameter: γ over attenuation at this frequency times the shock formation distance. Notice that the Burgers equation has only one free parameter associated with it and that is this parameter, γ . You have this in the review articles; it is referred to as the Goldberg number.

What this does is tell you the relative importance of dissipation versus nonlinearity and completely parameterizes the extent to which this plane wave will or will not shock or attenuate in a classical thermoviscous fluid.

How do we solve this equation? Assume a Fourier series expansion over these anharmonic components. Substitute into the Burgers equation. Once you have done that, you can eliminate time. What you have got now is an evolution equation for the rate of change of every one of the spectral amplitudes in terms of the attenuation at that frequency and in terms of nonlinearity, which in the frequency domain comes out as a discrete convolution here.

The advantage of this is now, first of all, this equation can be solved very rapidly by simple techniques such as Runge-Kutta integration methods -- it is very fast. Secondly, now that you are in the frequency domain, you are free to basically stuff in any type of attenuation or dispersion law you may wish.

For example, this parameter can be exchanged for some *ad hoc* real and imaginary part that would give you an arbitrary attenuation and dispersion law, for example, for relaxation in fluids, for boundary layer attenuation in ducts, or what-have-you.

(Transparency)

I actually gave you the lines of code here just to show you how really trivial it is to solve that equation. The top lines here are just packing the arrays. These four lines are your entire Runge-Kutta algorithm. The attenuation parameter here is a complex number that gives you arbitrary attenuation and dispersion -- you can read this out of an array. The sum terms are these discrete convolutions. That entire algorithm is just this sub-routine.

You can type this thing in, in 10 minutes. You can do some very interesting problems with this. You are free to choose whatever kinds of source waveforms you want. If you want to do pulses, you can just repeat your waveforms periodically in time, such that you are starting with a periodic waveform. You can use any dispersion-attenuation laws you want. You can include spreading in the wave by putting in simple geometric factors, for example, to account for spherical spreading or cylindrical spreading. It is a very convenient and simple tool.

(Transparency)

What does the output from this look like? Here is a typical output. Again, this is not done to get high precision, but just to show you some of the numerical aspects. This is 75 harmonics used to distort a sinusoid. The Goldberg number is 100. That means nonlinearity is strong compared to thermoviscous dissipation.

Here is the wave at the shock formation distance. When you go to two shock formation distances, because we have not used, for example, enough harmonics, you will find Gibbs oscillations here. You go through and count and you get about 75 of them. If we were to double the number of harmonics, these would virtually disappear. We get our proper sawtooth solution coming out of this algorithm.

This is an alternative method. It is simpler to code, but it has its own restrictions to the algorithm I showed you that was developed by Pestorius and Blackstock.

(Transparency)

I am going to skip over sound beams -- I am going to put that at the end and maybe hit a few highlights of that. Before the break I wanted to start to continue on with the discussion about dissipation. I wanted to say a few words about relaxation, because of its relation to this algorithm and because of its relation in the following topics to ones that came earlier. So skip over sound beams and move up to the relaxation just to show you how this additional mechanism can be included in the equation.

(Transparency)

This is another mechanism that could be included, according to the equations that were presented to you earlier. What we have used here is really the phenomenological approach from Landau's book. We do not concern ourselves as to what, exactly, is the relaxation process; all we are saying is we know that there is a time associated with going to a new equilibrium state.

For example, it could be molecular modes of vibration, phase transitions, chemical reactions, what-have-you. The main point is that these transitions are described by a new thermodynamic quantity, which he labels C , the rate of change of which is proportional to its displacement, whatever that may be, from equilibrium.

If there is a step jump in one of the thermodynamic variables, there is a characteristic relaxation time, τ , associated with that particular process. If you assume a generic description like that, you can go back to your equation of state and modify it as follows.

Here was our linear term in our pressure-density relation, our nonlinear term, quadratic, and what you do is you pick up this convolution integral that describes the relaxation process (this time variable should have a prime, because it is being integrated over), and M is what is called the dispersion.

C_0 is the sound speed at very low frequencies. C_∞ is the so-called frozen sound speed that we have heard about at very high frequencies. This equation can be transformed into either the

frequency domain to recover, after it is included with the wave equation, the attenuation and dispersion laws that we saw earlier, or it can be incorporated into the model equations that we just derived in an equally straightforward way.

I will just show you what it does to the wave. For a single relaxing mechanism, if we separate the attenuation from the dispersion, we see attenuation increases as frequency squared, where we are much less than the relaxation frequency of the medium, 1 over the relaxation time. In this case, one argues that perhaps it could be lumped in with bulk viscosity; it is indistinguishable, certainly, in terms of its frequency properties, and it plateaus off at the high frequencies.

The sound speed, the phase speed, increases from the equilibrium phase speed, C_0 , at low frequencies through transition at the relaxation frequency to the frozen sound speed, C infinity, at the higher frequencies. You can include additional relaxation mechanisms. All it will do is give you additional plateaus.

For example, in air to a first approximation you might want two, one for oxygen, one for nitrogen. You might want two in water for the various salt association relaxations. But you can, just by superposition, add additional mechanisms in, superposition because the relaxation is a second-order term, so it is a loss term times a linear term. Therefore, those loss mechanisms can be added, exactly.

DR. GARRETT: There is a very subtle and yet very powerful thing that you have slipped in here. If you define the M as the dispersion, and that shows up in your draft, Figure 5.1, $M/2$, this is the square root, really. You are assuming it is small, right?

DR. HAMILTON: Right.

DR. GARRETT: It also shows up in your equation of state, which is the interesting thing. That is the same M again, isn't it? Coefficient of the integral term?

DR. HAMILTON: Yes.

DR. GARRETT: That is really interesting. Is that Kramers-Kronig or something subtle, that you know that those things are correct?

DR. HAMILTON: This relaxation law, at least to this approximation, satisfies the Kramers-Kronig relations.

DR. GARRETT: Right. The wonderful thing is that the M shows up in both equations.

DR. HAMILTON: Both this equation and one other one?

DR. GARRETT: Exactly, in its defining equation below.

DR. HAMILTON: I am not sure what the question is. M is a very, very small parameter, a part in a thousand; for example, frequently.

DR. GARRETT: No, it is an observation. I was hoping that there was some obvious mathematical connection that would show that the M in the expansion of the equation of state was

obviously the M as you have defined it, and I think it is a good definition, for the dispersion strength. M is dispersion strength, dimensionless.

DR. HAMILTON: Exactly.

DR. GARRETT: I just wanted to know if there is an obvious way to show that those were, in fact, the same M 's.

DR. HAMILTON: Those are the same M 's.

DR. GARRETT: I know they are. Is there an obvious way to show that they are the same M 's? I understand the definition of M , but I do not understand --

DR. HAMILTON: Why does it pop up in the integral?

DR. GARRETT: Yes.

DR. HAMILTON: Why not?

(Laughter)

DR. GARRETT: Well, because this M is tied to that τ in this integral, and I find that fascinating. I do not say it is right or wrong or anything. It is a piece of physical acoustics that I suspect is true, but I do not know. We can take it off line, because that is really esoteric.

DR. MARSTON: I think there may be another way to look at it in terms of the Kramers-Kronig relations. The τ specifies the location of the pole if you think of a dispersion relationship in the complex frequency domain -- for those of you from E.M. background, the complex dielectric function. The τ specifies the location of the pole and the M specifies the residue or the strength, so there are two parameters that show up and that is all it is.

DR. GARRETT: Thank you.

DR. HAMILTON: That will work.

STUDENT: Can you give me the definition of relaxation?

DR. HAMILTON: The definition that I am using here is completely phenomenological. It is just the assumption that there is a time-dependent factor to transitions in state in the fluid, and I do not know what they are. Perhaps it is energy exchange with the molecular mode of -- for example, two chemicals in mixture. As I change the pressure, it will take a while to reestablish a new equilibrium state. That may be a huge relaxation time, but there is a time associated with it, nonetheless.

DR. SWIFT: In thermoacoustic stacks the channel widths are comparable in penetration. You might have four viscous.... In that case, can we start with the Burgers or Westervelt or is the approximation that the weakness of the attenuation lets you throw things away because it is low order?

DR. HAMILTON: We have thrown out vorticity, also, and vorticity would be important in your stacks, because of the shearing, which would give you the streaming. So no, you would have to start over.

DR. GARRETT: You said that by introducing this relaxation you can lump it into this thing that is called bulk viscosity.

DR. HAMILTON: At low frequencies.

DR. GARRETT: Yes, well, I would claim that that is -- you are not lumping it into it, that is the bulk viscosity. There is no other bulk viscosity.

DR. HAMILTON: Until I came to the summer school, that is what I would have said. I am being careful with my words here.

(Laughter)

Now that you have brought up Kramers-Kronig, let me add one point here. Kramers-Kronig, it is a consequence of causality. You get it through Fourier transform theory. If you know the dispersion of the system, you automatically know the attenuation, you know the complementary real and imaginary parts of these variables.

The relaxation does satisfy the Kramers-Kronig relation. Sometimes what people do is -- for example, in biological media they approximate the attenuation in tissue as attenuation varying linearly with frequency as opposed to quadratically in a thermoviscous fluid.

The question is, does that satisfy Kramers-Kronig? No, it does not, so the question is, should there be dispersion laws to be coupled with that? One way to check these conditions is you can build attenuation laws by adding up relaxation mechanisms theoretically. You know each of them individually satisfies Kramers-Kronig, so you can construct these different attenuation laws and know that you are satisfying the basic equations, yet you get the attenuation properties you want and then find out if dispersion is or is not an important effect there.

(Transparency)

Let me show you what relaxation does to a sound wave. In fact, this was calculated just with that spectral code that I showed you a few viewgraphs earlier. The source frequency is roughly comparable to the relaxation frequency; it was just doctored to amplify the results. What you see is, instead of getting a symmetric sawtooth, what you get is, first of all, a phase advance of the wave, the shock is occurring earlier in time. That is due to the dispersion associated with the relaxation.

Secondly, you see that after the shock hits, there is a kind of a rounding of the waveform that is associated with the relaxation at this point. You see the same thing in sonic booms because of the relaxation of the atmosphere. As the leading edge of that shock hits, you get a rounding of the wave profile immediately following that shock, primarily due to the effects of relaxation.

The example I just showed you of relaxation is an example of weak dispersion, so it is included in the equations of motion just like the viscosity, just like the heat conduction. In this case, the small parameter is this dispersion M .

(Transparency)

In the case of tubal attenuation for waveguides for the pipe experiments I described, shear viscosity at the walls and heat conduction give you some dispersion. Again, this is a weak effect. The dispersion is weak in the sense that it does not substantially reshuffle the harmonics that are being generated in that wave during propagation.

In particular, the coherence of the spectral components is maintained over distances that are long compared to the shock formation distances, so that shocks can occur. This is the major distinction, really, in terms of interpretation of waveform distortion in nonlinear acoustics and nonlinear optics.

When you are looking for nonlinear optical phenomena, the materials in which the nonlinearity may be extraordinarily high also have very high dispersion. There are, however, a few acoustical systems in which you do get very large dispersion.

One of them is waveguides. In particular, for non-planar modes in waveguides these phase speeds can be arbitrarily large. As you approach cutoff, phase speeds are going infinite. So by tuning the frequency and varying the duct width, what-have-you, you can have arbitrarily large and different phase speeds between different interacting harmonic components.

(Transparency)

What I am going to show you is just one example of nonlinear interaction and harmonic generation in a waveguide. This is an experiment done at Texas, also. The JBL driver on the end launches a plane wave mode, the 0-0 mode in this rectangular duct, about this big. These drivers on the top and the bottom are attached to slits near the end of the waveguide in opposite phase and are driven in push-pull in such a way as to generate the first lowest order mode of this high-frequency wave propagating down the duct. These two wave speeds are going to be different, corresponding to the wave vectors of their angular decomposition.

So imagine what is going to happen here if you think back to our perturbation method. These are our linear solutions for the waveguide. We have two primary waves, a plane wave and a second wave in a higher order mode. They are going at different speeds, different frequencies.

At second order, these spectral components are going to interact. They will generate a sum frequency wave. They will also generate a difference frequency wave. The primary waves are now pumping energy into these sum-and-difference frequency waves at different speeds from which they would ordinarily want to propagate.

The spectral components that are being generated are going at speeds that are yet again different from those that are doing the interaction. What this is going to do is give you an asynchronous situation, a phase mismatch, which will temporarily pump energy into these harmonics and then suck it back out, again, pump it in, draw it out. It is like pumping a mass spring oscillator off resonance; what you have got is the particular solution, due to the forcing function, is not in resonance with the freely oscillating solution of the homogeneous solution; these

solutions will beat. Instead of beating in time, what we get is beating in space. That is exactly what happens.

(Transparency)

Here is the generation of the sum-frequency wave. Here is the generation of the difference-frequency wave. Indeed, these oscillate as they go down the waveguide. This is the effect that is discussed frequently in nonlinear optics, but measurements like this of these amplitudes are almost prohibitive in optics, but they are more accessible in nonlinear acoustics.

This demonstrates the phase coherence of these interactions. As the dispersion in the system decreases and these waves propagate more toward the same speed, the dispersion length we say increases and these waves have longer opportunities to amplify before they start coming back down, again.

The whole key here is are your dispersion lengths related to what are your shock formation lengths in terms of this nonlinearity important or not in a dispersive system?

(Transparency)

Another highly dispersive system that gets a lot of attention is bubbly liquids. These are extremely dispersive. Here I will give you some of the basic equations, especially with Dr. Atchley having derived the Rayleigh-Plesset equation this morning. The idea is this.

Bubbles are resonant oscillators. As the sound wave, the pressure disturbance, propagates through this mixture, it excites these bubbles, energy is exchanged with these bubbles, we have these time lags that we have talked about before, and that will give rise to dispersion. It will also give rise to large attenuation but the dispersion is what we will look at here initially.

How do you derive a model equation for this system? What we are going to assume is that all of the nonlinearity in the system is due to the bubble pulsations. As we saw this morning, these bubble pulsations are indeed highly nonlinear. We will ignore the nonlinearities we have talked about already, the B/A of the liquid, the convection, and so forth. All our nonlinearity will be due to the oscillators.

We can begin, then, with a kind of equation of state, if you will, for the mixture. ρ is the density of the mixture, ρ_0 is its equilibrium value, ρ is its disturbed value. V is the bubble volume, instantaneous, and N is the number of bubbles per unit volume in the liquid. All of this is in this reference by Zabolotskaya and Soluyan in 1973, so I will not go into too much more definition.

The upshot is you begin with a mixture rule for the densities, where you have the bubble volumes, the compressibility due to the bubbles, and the compressibility due to the liquid, that is, that which has been displaced by the bubbles.

You now put in your acoustic approximations. Let us expand the mixture density, let us expand the density of the liquid, and let us expand the density of the gas, do some tricks up here,

and the linearized form of that is this linear equation of state: excess density in terms of sound pressure with a correction term, where N is our bubble density and B is the excess volume of the bubble (in other words, V_0 is its equilibrium position and v is its excursion about that equilibrium position). So we have a linearized equation here now for the compressibility in the equation of state.

Now let us linearize the mass conservation equation, let us linearize the momentum equation, use this as our equation of state, and stick them together. We have an inhomogeneous wave equation, sound pressure on the left and a linear term in the excess bubble volume on the right. Here is our wave equation. We have got one equation and two unknowns: sound pressure, bubble volume.

We need a second equation to relate these two and this is where the nonlinearity kicks in, that is, through the Rayleigh-Plesset equation governing the bubble dynamics. This is the equation Dr. Atchley derived for you this morning. R is the radius of the bubble, P_g is the gas pressure inside, P the pressure outside related to the sound pressure that will be impinging upon this bubble.

What we do now is say let us not write this in terms of radius, let us write it in terms of velocity. Here is our bubble volume. Here is our relation to accomplish that. We are talking minor volume oscillations here, we are not talking sonoluminescence or anything like that. So let us assume a simple adiabatic gas law for the gas inside that bubble.

Combine this, combine this, and combine this, and write it to quadratic order in terms of the excess bubble volume, and this is your bubble oscillator equation. The nonlinear bubble terms are on the right-hand side. Let us ignore those for a second.

Look at the first three terms. The second derivative of the volume with respect to time, first derivative in the volume itself. This is a linear oscillator. ω_0 is the resonance frequency of this bubble oscillation (this is this Rayleigh frequency that has come up several times) and δ is a viscosity damping coefficient for this bubble.

P is the sound pressure exerted on this bubble. That is your forcing function. So we have got an harmonic oscillator that is driven by the sound pressure, just as we had this morning, with its own natural frequency.

DR. GARRETT: What is η ?

DR. HAMILTON: A combination of parameters. I forget what is inside it right now. A and B also have densities and bubbles mixed in with it, equilibrium bubble volumes, equilibrium bubble radii -- these are just algebraic expressions.

On the right-hand side you have got velocity squared. This is where the nonlinearity comes in. These are the nonlinear correction terms to the left-hand side. In this case, A is proportional to the adiabatic gas constant, so this is the nonlinearity in this equation of state. This nonlinearity is

coming from the Rayleigh-Plesset equation. These two can actually cancel, depending on how you are driving the system.

In any case, now we have two equations and two unknowns: the wave equation here in volume and pressure and the bubble equation here in volume and pressure. We can solve these by the same perturbation techniques I showed you in the beginning.

To linear order I throw out these terms. If I think of a single-frequency drive, I can solve for the bubble volume -- I get an omega squared here, omega here. I can invert this algebraic expression, put it back in up here, and I have a single expression in terms of the sound pressure going through a bubbly liquid.

From this I can extract my dispersion relation; in particular, the attenuation and the phase speed of the sound that propagates through that bubbly mixture.

(Transparency)

And here it is. This is what the attenuation and dispersion look like in a bubbly mixture. The bubble resonance frequency here is at unity; we have normalized by the natural frequency, the Rayleigh frequency of the bubble. Delta is a viscosity coefficient. Just think of it as some artificial damping that we are putting in for the time being.

At very low frequencies the bubbles are oscillating in phase with the sound wave; therefore, it simply looks like we have got a squishier fluid. The compressibility has increased, stiffness has gone down and, therefore, the wave speed has gone down. We get a slightly slower phase speed here than we do at much higher frequencies, in the high-frequency limit.

At very high frequencies the oscillation of the sound wave driving the bubble is happening too fast for the bubble to respond. The bubbles are effectively frozen and as the wave propagates through this mixture it does not even see the bubbles in the sense that it is responding now to the compressibility of the liquid, not the liquid-gas mixture.

Clearly, if you go to sufficiently high frequencies you are getting acoustic wavelengths on the order of the bubble size. We are not considering that limit here, where they become scatterers and so forth. So always the bubble sizes are assumed to be small compared to a wavelength.

In this resonance region things go crazy. If you do not have attenuation in the system, your sound speed goes from zero to infinity through a transition region and damping, of course, smooths that out. Attenuation is very large, also, around the bubble resonance frequency. Typically, again, in relation to Kramers-Kronig, you are not getting huge dispersion here without huge attenuation. That is one of the offsetting issues here.

People have looked at bubbly liquids as ways of doing parametric amplification of sound waves because of this extraordinary nonlinearity. Some of the problems involved with that are how do you get these uniform bubble mixtures and, in addition, now that you have got the

nonlinearity, you are stuck with the attenuation, too. These two effects could compete and offset what you are trying to do. That is the linear solution.

Now, let us take that linear sound-wave solution for the sound field and stick it back into the right-hand side and solve for the second harmonic field the way we might have for the plane waves earlier.

(Transparency)

You will get an equation and system looking like this. Linear operator on the left-hand side, a complex sound speed for the second harmonic frequency -- this complex sound speed is where your attenuation and dispersion come in. On the right-hand side, the driving function is proportional to the primary wave squared -- that is our omega wave coming through -- and a coefficient of nonlinearity that now, in essence, depends on the frequency of that primary wave.

This nonlinearity coefficient is plotted in the lower picture here. You can see that it becomes almost arbitrarily large without dissipation, if the dissipation is weak in the system. If the drive frequency is matching the bubble resonance frequency, that is one case. Secondly, you have a very strong amplification of the second harmonic wave if you drive the system at one-half of the bubble resonance frequency.

If you are at one-half the bubble resonance frequency, you are now coupling into the second harmonic and driving it at the bubble resonance frequency. In this case the primary wave is at resonance. Here the second harmonic is at resonance. Either way you get strong coupling and energy transfer to the second harmonic component.

You get a zero out here, and that is where the nonlinearity in the Rayleigh-Plesset equation cancels it. I spoke about the Rayleigh-Plesset nonlinearity cancels with the adiabatic gas law nonlinearity and you actually go through a zero out here.

(Transparency)

I think it was in yesterday's talk that Dr. Maynard mentioned that people often go out of their way to derive nonlinear Schrödinger equations. Here is one, a KdV equation, in this case.

(Transparency)

What I am going to assume now is that we are operating down here in the low-frequency portion of the wavefield, so the dispersion is relatively weak, so I can consider waves with frequencies substantially below resonance and make approximations from there.

(Transparency)

Just to hit the highlights, we start with the wave equation from the first slide. Here is our bubble volume dynamics equation after we have made the assumption of low frequencies. Next, because all of these are small terms, I can use this relation between pressure and bubble volume to eliminate volume here, here, and here, so now I have volume in terms of pressure across the board.

Take this, put it in up here, I have an equation now just in terms of the sound pressure, P . The next thing I do is say, okay, I am down at low frequencies, let me put myself in a moving frame of reference, as when I derived the Burgers equation at the low frequency sound speed, C_{00} -- that is that low-frequency limit of that dispersion curve that I showed you.

Secondly, let us assume, again, slow variation with distance, as we did for the derivation of Burgers equation, make these substitutions, put them in up here, and what we end up with is something referred to as the KdV-Burgers equation.

If I eliminate the third term here, this is what is called the KdV equation. Here is our quadratic nonlinearity in the system due to the bubbles. This, with three time derivatives, linear term, gives me dispersion. If I eliminate the B term, these two derivatives correspond to a kind of viscosity damping in the sound wave -- here is the quadratic nonlinearity -- and I get the Burgers equation. Including the two, the name is picked up of the KdV-Burgers equation.

The interest here is that if the dissipation is very small, people look for soliton solutions, as you have heard about. Let us exclude the dissipation and come back to this equation, which is the KdV equation.

(Transparency)

The excitement here is that this type of equation describes, now, two physical effects: the tendency of a waveform to steepen due to nonlinearity and, second, there is a dispersion in the wave. There is no energy loss in this model, only dispersion.

The question is, is there a particular type of pulse shape for which the dispersion just perfectly offsets the tendency of the wave to steepen forward? When that exists for a single pulse, this is called a soliton, or a solitary wave solution. Many such solutions of this equation exist. The case of a single soliton is given by this -- it is just a blip like this. It has a characteristic velocity, excess velocity, V , compared to our equilibrium sound speed. Second, it has a pulse duration, T , associated with its width.

What these equations tell you is that the higher the amplitude of the pulse, the faster it goes, and the higher the amplitude, the narrower it must be to be a solution to that equation. So these solitons have the form that tall thin ones move fast and short fat ones move slower. These will catch up with these.

Solitons can actually pass through each other. If you let these waves propagate, eventually they will combine. Keep in mind that super position does not hold, because this is still a nonlinear system, but they will pass through each other and they will look the same.

What is somewhat interesting is you can actually know if a collision took place, because if you plotted these two, for example, as a function in space and time and you had two solitons that were about to interact, the faster one is the tall thin one, the slower one is the shorter fat one.

If this interaction were linear, you would expect their trajectories to come out the way they went in. They do not. I do not remember which way they are displaced, but they come out displaced. So the pulses do, in fact, interact, but you do get exactly the same pulse shapes coming out the other side. It is not just a pure passage between them; they do have a particle-like behavior associated with them.

(Transparency)

Have these been measured? The only paper I can find is this one from some Russian scientists in Novosibirsk. At any rate, I give you that just to have the reference for this particular article. They claim to have found soliton-type solutions.

These equations also emit multiple soliton solutions, so you get, actually, families of solitons and watch them collide. That is what they are looking at here, for the first case.

STUDENT: so the nonlinear items appear after the bubbles

DR. HAMILTON: The nonlinearity is weak, yes, and the second part of your question?

STUDENT:

DR. HAMILTON: I am not sure if I understand your question, but we are not including any scattering of sound from the bubble. We are considering this just as a mixture, as a uniform homogeneous mixture.

STUDENT: But the amplitude and the width are small?

DR. HAMILTON: Yes, relatively small in the sense of small acoustic Mach number and weak oscillations of the bubbles. Yes, it is a weak nonlinearity limit.

MR. HALLEJ: The solitons that you drew there on the board, you say that they are displaced. I am having a hard time picturing that for a case where, let us say, the trajectories of them were at 180 degrees or zero degrees. Let us say the tall thin one comes up and catches up along the same line. I mean, how do they displace? It does not seem like something you would expect to see. Or if they are coming head on at each other.

DR. HAMILTON: No, you do not, but what you are saying you expect to see is based on linear superposition, that one passes through the other without slowing down or speeding up.

MR. HALLEJ: Yes, the case you drew is kind of the harder case to imagine, but I am wondering what it would be like if they come right along a straight line at each other. How does this displacement happen?

DR. HAMILTON: While they are interacting, they are no longer the same pulse; it is a different entity. That may be a much smaller change, because their interaction time is much less, I would presume. But I would expect to also see a displacement in the space-time frame. In fact, I would be surprised if you did not.

STUDENT: You just said that the narrow one moves faster. Is that generally true?

DR. HAMILTON: I believe solitons that are solutions of the nonlinear Schrödinger equation have the same property. The higher amplitudes are going to move faster. In those cases, the solitons are called envelope solitons. You have a tone burst and it is the envelope you are now following that does not change in shape.

I think you would have to look system by system. I do not know of a general rule for that, but for the KdV, yes.

(Transparency)

What about dc effects? I will say a few words on this. If we now do our expansion and go to our second-order quantities, let us not look for the second harmonic any more; let us take a look at the zero frequency term. If we are looking at particle velocity, this corresponds to flow and we call it acoustical streaming. If it corresponds to dc pressure, we call it radiation pressure, and that gives you, for example, levitation that we talked about earlier.

We will use expansions and do successive approximations the same way as before, where the dc pressure and these Eulerian coordinates are simply the time average of the Eulerian particle velocity, rather, over, say, a period or multiple periods of that waveform.

The only comment I will make here in terms of analysis is you have to be careful, when you are reading the literature on acoustical streaming, because there are about three streaming variables that kick around. One is the Eulerian velocity -- that is the one you might measure with a hot wire anemometer. You are at a fixed location, you are looking at flow moving by a point.

There is, then, the Lagrangian dc velocity, so you are following the particles. Oddly enough, there is a third velocity that is mathematically different from the Lagrangian velocity and that is called the mass flow-rate velocity, where you force a quadratic order continuity of mass.

So you have these three possibilities for the streaming velocity that are taken into account. Oftentimes their differences from each other are trivial, but the equations are radically different and you can wonder what is going on when you take a look at one versus the other.

(Transparency)

What do the driving equations look like? The best review article I know on this by far is Wes Nyborg's article, called "Acoustic Streaming," in the series called *Physical Acoustics* by Mason back in 1965, an outstanding review of this whole field at that time.

(Transparency)

What we do is we go up into the mass and momentum equations, combine them, time average, and what we do, then, is we find the dc velocity as a function of the dc pressure and the radiation pressure plus a driving term, forcing function, that is a function of the primary wavefield U_1 or P_1 .

Usually the dc pressure in these analyses, particularly for progressive waves, can be thrown out, so what you have got is simply a kind of Laplace's equation here: dc velocity is dependent on

this driving force. In the case of plane waves, the force driving the streaming is very simple; it is the gradient of the particle velocity squared, the X derivative of the particle velocity squared.

For example, for a plane wave, frequency ω , attenuation α , it is a very simple solution for the force acting on the fluid. Why does this streaming occur? It is because of the attenuation. As the wave propagates, it is being attenuated, it is losing momentum. Where does that momentum go? It is going into pushing the fluid in dc flow.

If you have a collimated beam of sound, a wave is propagating this way and attenuating, it will drive the dc flow in this direction. Mass must be conserved. At some point you are going to have to have some kind of circulation so that if you draw a plane and integrate across it, you will have zero mass flux going through it.

(Transparency)

This is what it looks like. Here is a source radiating into water, an ultrasound source, and, indeed, here is the streaming going along the axis and circulating back (it is evidently in a container here, bound above and below) along like this.

To get back to bulk viscosity, this streaming is actually used to measure bulk viscosity, whatever that is. Yes?

MR. PRATHER: On the previous picture you had the return on the outside of the particle.

DR. HAMILTON: These are not mine. It is just meant to mean the edges of the sound beam. So no, we are in free space here.

(Transparency)

Check this out. Look at your equation that we are solving on the top slide here -- forget about the pressure. The velocity we are solving for is multiplied by the shear viscosity μ .

Now let us go to the right side. The forcing function is proportional to α , but α is proportional to the shear viscosity μ -- μ cancels on both sides of the equation.

On the one hand, you need attenuation to get the streaming. On the other hand, when you are done, it is not even in the equation any more. Kind of strange.

If you have bulk viscosity, that does appear explicitly. As you can see, this has to do with the basic nature of the equations. In most of the quasi-plane results for progressive beams, the shear viscosity coefficient has dropped out.

(Transparency)

Eckart derived an analytic solution for the following simple case, a piston in a tube. He said, okay, let us let the sound beam be high KA, very collimated, and let us put a permeable membrane here, like a Glad bag. It is a low-frequency filter. The sound wave does not see it, a high-pass filter. The sound wave goes through it, but it is going to trap the dc flow. So any flow is just going to circulate in this cavity.

Now let us look at a little region in-between and you can actually do the analysis now based on these simple plane wave equations for the streaming across the duct. You get a very simple analytic result for this. You find that in the middle of the beam, sure enough, the streaming is going out and it has got to come back along the sides and you get expressions for the streaming velocity as a function of piston radius versus a duct radius.

The only reason I wrote this slide is to show you that the peak streaming velocity on axis is given by this expression and is proportional to the K of the source, particle velocity of the source, bulk viscosity over shear viscosity.

This was used by Hank Medwin when he was a graduate student -- he is at the Naval Postgraduate School, or has been. You can use this, if you know the shear viscosity by another mechanism, you can measure the bulk viscosity, and that is what he did for his dissertation. It works.

I will skip the next slide, streaming and sound beams.

(Transparency)

Standing waves become more tricky. Now you have got shearing. Vorticity is coming from the side walls. The analysis in this case is extremely complicated algebraically. This is the case of a standing wave in a duct that I have shown. Your pressure oscillations are going like this for your sound field.

Now you have these circulation patterns inside the duct. It is more complicated. You have large-scale outer streaming vortices and smaller-scale inner streaming vortices and these things are just going like gears meshing in with each other. These outer-streaming vortices were first derived by Rayleigh.

(Transparency)

And photographed by someone named Andrade in the early part of this century. This is a smoke-filled tube and he shows precisely these streaming vortices. What he could not see is the inner streaming vortices near the boundary layers.

(Transparency)

But to show they exist, this is a neat picture out of Van Dyke's book of photographs in fluid mechanics (I forget the name of it). There is a mathematically analogous problem for an oscillating cylinder in an infinite fluid, and here you see the inner streaming vortices. Here is the cylinder, it is oscillating along this axis. Here is the outer streaming pattern -- here.

In this case, the streaming can be so severe that the inner streaming vortices produce a jet, which can fire out along the axes of the cylinder. This is a rather well-studied classical acoustical streaming problem. This is done in water and glycerine.

DR. GARRETT: Is there a jet out in the horizontal direction....vertical directions?

DR. HAMILTON: Yes, the jet goes out in this way and it is feeding in this way. The engine that is pumping that jet are not the outer ones. It is the inner patterns that are doing the ejection.

(Transparency)

Finally, some comments about radiation pressure. Suppose you have got a sound wave in a standing wave tube and you go back and you time-average your momentum equation. You can come up with two types -- there are actually four names of radiation pressures that kick around in the literature -- but there are two that come immediately.

One is from the equations we derived, a so-called Eulerian pressure, because we are in Eulerian coordinates when we did the averaging: It turns out to be this Lagrangian density, plus a constant. We can evaluate this for the standing wave in the pipe and see that it oscillates like this.

At pressure anti-nodes the Eulerian pressure is a maximum. At pressure nodes, particle velocity minima, the pressure here is a minimum. So this is the dc pressure, the radiation pressure that you would expect to see in this tube.

If you were now to go into Lagrangian coordinates and ask what is the Lagrangian pressure, radiation pressure, it is a constant; it does not even depend on X . What do we mean by Lagrangian radiation pressure? It means you are riding with a fluid particle while you are taking your time average. The only places these two expressions match up are where there are velocity nodes, where the fluid is not moving.

The one we are interested in here is the Eulerian pressure in the tube. What does this look like?

(Transparency)

One way to visualize this is to put water in the bottom of a standing wave tube laid on its side. Remember, at the ends we had radiation pressure maxima, so this will pressure the water down. At radiation pressure minima we will expect it to arch. We have got about a 150 or 160 dB wave in the sound field here (this is air above, water below). These waves are not moving; this water just arches up like that.

If you calculate millimeters of displacement and calculate hydrostatic pressure, it is about what you would expect from these calculations for the Eulerian radiation pressure.

This was done by a former summer school student, Tom Van Dorn, as a class project to illustrate this. Like Felipe, he kept turning the knob up, and this is what happened.

(Transparency)

Things started fountaining like crazy. Here it was with 160 dB. Here he has another half-dB. The surface of the water is arching a little bit, it starts getting a little capillary action and it is feeling a little jittery and unstable, and then eventually this violent fountain starts here, and fountains occur at every anti-node in the displacement of this water surface, and these things will go forever.

What we think is happening here -- we do not know how the instability starts due to the capillary wave action, or what-have-you -- but it seems this inner streaming vortex, which is in the proper direction at that surface for ejecting this fluid. You will see this as a demonstration down at the Exploratorium, as a matter of fact, as one of their exhibits.

MR. HALLEJ: About what frequency is that?

DR. HAMILTON: Several kilohertz in air.

(Transparency)

Stranger things happen in these standing waves. If you take the water out of there and you put a powder down and look at the patterns you get, a half wavelength is this, so this pattern is repeating itself. The sound wave causes a striation of the pattern along the lower wall of these tubes, and these patterns are just remarkable. I mean, they are like razor-sharp, lined up like soldiers there. It happens very fast and if you look at the upper picture here, these little towers of particles are very, very thin, often single particle in width.

The sound wave somehow gets down to this scale, an order of magnitude scale less than that of the wavelength, and somehow is able to align these various particles -- this is more of a curiosity than anything else. This is also from Andrade's paper. He has a number of experiments on that. Just a passing fancy here.

DR. ATCHLEY: If those were bubbles, you would say they were Bjorkes forces. Is that seen here, does anybody know? Rayleigh also had an explanation, but is it Bjorkes scattering, do you know, Phil?

DR. MARSTON: No, I do not know, but the closest thing that has been analyzed is Mary Ann Holtewanger-Weiser's work at Yale, where she found that in addition to the normal Bjorkes force layering of particles, you can get sublayering in between, but I do not believe she related it to this work, and it is unclear to me whether they are the same thing.

DR. GARRETT: This effect was explained in Rayleigh's theory of sound in a very simple way. I can do it after we break. Basically, it is just the Bernoulli force. I tend to call it the "Maidenform effect," because it both lifts and separates.

If you two particles like this, then the fluid flow through them is faster and the Bernoulli pressure brings them together in a vertical direction. But in the vertical direction, then, they form a barrier and you get stagnation in between, and that causes the striations. It is covered in Rayleigh Volume II.

This was done very early in Kundt tubes, because they could stroke a rod and have high-amplitude sound and get a standing wave. The standard means for measuring wavelength was to throw

DR. HAMILTON: But were they getting striations?

DR. GARRETT: Absolutely, striations, and they were getting the lifting kind of Bernoulli forces and the stagnation gives you the striations. I do not believe that Rayleigh had a quantitative theory for the actual periodicity of the striations, but he absolutely -- there are beautiful pictures. In fact, it is in my notes from two years ago.

DR. KEOLIAN: Although it may go with the particle size, because they used to demonstrate this in his levitation tube with styrofoam peanuts and they would line up in sheets and they would separate by peanut size, whereas these are much thinner.

DR. GARRETT: I will guarantee that Rayleigh does not have a scale length that tells you what that separation is going to be *a priori* from the particles or anything else.

DR. SWIFT: In our group Jeff Olsen tried to make some measurements of the cosine of TKX pressure variation, not with water, but putting in a little tiny capillary with a tiny, tiny hole and pressure transducer and this was 30 bar helium pressure oscillations in 5 to 10 percent. It was better than a factor of two, but not in perfect agreement with what you would expect.

DR. HAMILTON: Better than a factor of two with respect to what? Amplitude?

DR. SWIFT:

DR. HAMILTON: So the amplitude was way off.

DR. SWIFT: static pressure oscillation as a function of position. Do you remember if Van Dorn's measurements were accurate enough to --

DR. HAMILTON: No, and we did not do it for that purpose. It was just for illustration. It was not research.

DR. GARRETT: There were experiments that Izzy did using a pith ball on a pendulum, so gravity and that was repeated by Varmass and Emily Leon in JASA, probably about 15 years ago. For the acoustics levitation guys at JPL it is important to know

DR. HAMILTON: This analysis is Taylor Wang's out of the JPL group, because there has been argument in the literature -- they are very subtle points -- about how you set the boundary points on these problems.

For example, the difference between Eulerian and Lagrangian: All of a sudden one effect disappears. So you have got to be very sensitive to how you set the problem up. This is one of the most recent papers out of that group.

We are out of time and we will not get to sound beams. Maybe next time. Thanks a lot.

ATMOSPHERIC ACOUSTICS

RICHARD RASPET
UNIVERSITY OF MISSISSIPPI

DR. RASPET: Since this is California, I have to put this warning up here, a quote from Allen Pierce's book: "A deep understanding of acoustics is not acquired with a superficial effort." This is a big warning, when you are taking a class out of that book and you open it up and read that.

(Transparency 1)

I want to give some background. Here is an outline. We are really not going to do all these things. The first thing I want to try to do in section I is to give you an understanding of what affects sound propagation in the atmosphere. My primary interest is mainly surface to surface propagation. I will not discuss such things as atomic bomb explosions over long distances.

We live close to the ground, and so what I will talk about things that affect people in a direct way. We are going to look at atmospheric refraction and the effects of the ground with a simple model, first. Then we will take a break. We have a laboratory on ray tracing, and that is why this equipment has been handed out to you. I am going to make sure that you will take something away from my talk, even if it is a compass, triangle, and ruler.

In Section II we will look at the right way to do these solutions, or what I think is the right way. The graduate talk involves -- if you are interested in these subjects, we will get together after this talk -- these sorts of things: Density variation is something that has been misunderstood many times over the years; winds are pretty tricky but we will give you an approximation that works pretty well; and the 3-D FFP is a calculation a graduate student of mine worked on.

What do these numbers in the margins mean? If you understand Section I, then you can be a consultant and you can work noise problems and earn hundreds of thousands of dollars. If you study more and learn the full-wave methods of Section II, you can be a professor and earn tens of thousands of dollars. If you are interested in the graduate subjects, you can be a graduate student and work hundreds of hours a week for tens of dollars.

(Laughter)

The first thing we should do is make sure you can be a consultant.

(Transparency 2)

Why are we interested in sound propagation? I put this in reverse order of where the funding comes from: Understanding our world. We communicate with sound, and so do many creatures. This is the "life, death, and pursuit of happiness" section of the talk.

(Slide of a bush cricket)

This little creature calls for mates. This is a bush cricket. This bush cricket starts out in the evening and, depending on the atmospheric conditions, tunes his burrow to increase his output 20 dB. He changes the shape of his burrow to keep his output high, attract more mates, and be more successful in life.

(Slide of camel cricket on tree)

Another interesting that Tim Forrest and I have been working on is showing that camel crickets optimize their height on the side of a tree for maximum calling area. If the ground surface is different, the optimum height is different, and the crickets call from a different height.

Tim has been gathering data for some massive paper. We have about five or six creatures that either fly or call from pretty much what you would calculate as the optimum height. One that calls from the ground covers about 16 times less area, 1/16th of the area, as one that calls from the optimum position.

(Slide of an African elephant)

Here is a real expert in long-range sound propagation. Does anybody here want to imitate an elephant call for me? All right, Professor Bass.

(No response)

Did you hear that?

(Laughter)

Does anybody else want to make an elephant call? Go ahead, Dr. Atchley.

(No response)

He has got it right, too.

Elephant trumpeting on television is a short-range call. This is the African elephant. You get, say, 10 elephants over about a 10-kilometer square. They are social creatures. The females come into heat for about a week every two years. The males are in musth, or receptive, about one day every two months. They are over this huge area.

Elephants forage over huge areas and then come and meet in the evening at a watering hole. How do they do this? They flutter their sinuses and produce infrasound at 15 Hz. That is why Professor Bass produced a sub-audible 15-Hz sound. Some of you may have felt queasy, but you could not hear it.

(Laughter)

In fact, it was health investigations of elephant keepers that revealed that elephants communicate infrasonically.

We will see much later in this talk that during the day there are meteorological windows in which they can communicate up to 10 kilometers infrasonically.

(Slide of a tornado)

Tornadoes are a strong noise source. Professor Bass has developed a detector for tornadoes. Here is a noise source we need to understand and avoid. We also need to understand the propagation of warning signals and Civil Defense sirens. A lot of work has been done in acoustic propagation on sirens. Izzy Rudnick, for example, worked on sirens and sound propagation.

I am not sure who he was funded by, but I am aware, in the literature, that he did work on sirens. It might have been Civil Defense. If you may have noticed there are sirens around nuclear reactors and they are all mounted in the same way.

(Slide of an artillery piece)

Artillery is also a noise source. If you want to predict how annoyed people are going to be around Army installations, you need to be able to predict how impulsive noise propagates.

The noises from guns are finite waves close to the source. Weapons systems can be located by their noise emissions. This has been a source of funding for acoustic propagation research every war. That tells you how often this research received funding.

Outdoor sound propagation research is typically solution-oriented. Little basic research is performed. Somebody is annoyed by a noise and he wants a quick answer. Someone wants to find a weapon and he wants some quick answers. It tends to be cyclical. If you have a problem, somebody spends money. Often very good minds are brought in for short term studies. One of Einstein's graduate students in gravitational theory worked in acoustic propagation for two years near the end of WWII and produced some beautiful papers. Due to the cyclical nature of this research you can publish the same paper every 20 years. It is also low-level research, where you do not find large groups and concentrated efforts, except for occasional issues.

Many times the researcher only completes the first section of my outline. The investigators get an indication of the relevant mechanisms, they declare victory and quit. But it is not the full answer. Hopefully, I will get far enough along to give you some semi-full answers to questions, and talk about wave propagation.

(Transparency 3)

The mechanisms affecting sound propagation are listed up here. Professor Bass talked about spherical spreading. If you have a homogeneous source, isotropic in all directions, the energy spreads out as $1/r^2$, or as the area, and the pressure is the square root of the energy and spreads out as $1/r$. I am going to be sloppy about this today, and I hope you students will stop me when I throw up power laws for r and do not define them clearly.

DR. GARDNER: Why do you call that excess attenuation?

DR. RASPET: I am not to that point yet. I have not talked about atmospheric attenuation. I will answer your question shortly.

Professor Bass has talked about atmospheric attenuation. These are mechanisms that we are very sure of. If propagation occurred with only spherical spreading and atmospheric attenuation we could calculate losses perfectly. In addition, if we plot this level it would fall off very rapidly. If you double the distance, the level decreases as $1/4$.

We want to emphasize the effects other than spherical spreading and atmospheric attenuation so we add these back in, and it gives more detailed plots. If spherical spreading and excess attenuation were the only mechanisms present the excess attenuation plot would be zero at 1 meter, zero at 100 meters, zero at 1000 meters. We use excess attenuation to describe all other effects combined.

We are mainly going to talk about refraction, that is, the bending of sound waves by velocity gradients, sound speed gradients and wind gradients. We avoid talking about diffraction in the first part of the lecture, but will return to it in section II. In brief, diffraction is whatever is not described by refraction.

The next effect is the surface boundary conditions: What is the effect of the ground?

We are not going to talk about turbulence and non-linear effects even though they are hot topics in atmospheric acoustics. Scattering from turbulence is a very large effect. We will get to a plot where I will show you how bad our predictions are if we do not include scattering. Mark Hamilton had lectured about nonlinear effects. They are important for high-level sources. Mark Kelly is working on scattering from turbulence, and he will can you more about it.

(Transparency 4)

The things that make atmospheric sound propagation particularly interesting are that the wavelength is long -- we will have diffraction effects -- the nonlinear effects occur with relatively modest levels, and ground surface effects. Some pretty peculiar waves occur, as we will see.

The wind is anisotropic. We are going to operate at the kindergarten level, as far as wind goes. I will discuss this more with anyone who is interested in it.

Another subject is the meteorological variability. You cannot go out and expect to have sound levels be stable for periods of 15 to 20 minutes. It makes the measurements difficult. You cannot go out with one sound level meter, walk away from a source, and make measurements. The meteorology is constantly changing. We will see what sort of variations are expected.

Last is this question of turbulence, which introduces the idea of weak scattering, which we are not going to talk about.

(Transparency 5a)

My talk got shorter when Dr. Garrett and others supplied us with the wave equation. I just want to point out that one solution to this is plane waves. This is the picture we will examine in Section I.

(Transparency 6)

A wave is going in a direction \mathbf{k} , $\mathbf{k} \cdot \mathbf{r}$. We get the velocity from the pressure from the impedance relationship. This is the unit vector in the direction of propagation; $p/\rho c$ gives you velocity.

(Transparency 7)

We want to look at plane wave reflection from the ground. The way this is normally done for outdoor propagation is in terms of the specific acoustic impedance, the ratio of the pressure to the normal velocity into the surface. translational invariance holds for a uniform surface with incident plane waves. The ratio of pressure to normal velocity at any point along the surface must be the same, and this can be expressed in terms of impedance. The normalized impedance given by the impedance of air, ρc of air, divided into the impedance of that surface.

If the surface is not changing too rapidly and most of the sound is refracted at normal incidence into the ground, the plane wave reflection boundary condition can be used even when a uniform plane wave covering all space does not exist. The key phrase for this that appears in the literature is "normal ground impedance meets."

(Transparency 8)

There is a very rich history of ground impedance research in the literature. I have listed some authors here. Delany and Bazeley were the first to realize that the measurements they made on acoustic tile -- everybody at Asilomar thinks our acoustics group is working on making building interiors sound better -- could be applied to the outdoor ground surface.

That is a semi-empirical model. Other people have used physical calculations of sound propagation in tubes as a basis for calculation. The ground is grains of sand or grains of quartz that do not vibrate, and gaps. How big those gaps are tells us how acoustically soft the ground is. People have used Rayleigh's solution or Kirchhoff's solution for sound propagation in pores.

Small pore approximations of these theories are reasonably independent of pore shape and a theory with very few parameters can describe almost any ground surface.

Some distinguished names are listed here. Sabatier's doctoral thesis is on the effects of seismic coupling, where the grains of sand move in response to the sound wave. This is a very low-frequency effect, and, for purposes of this lecture, can be ignored. All we take into account are viscous and thermal losses in tortuous tubes, tubes that twist, and we can characterize the ground with a few parameters for how porous the ground is. The more porous it is, the more sound it lets in. The flow resistance is one parameter that relates to how "twisty" the tubes are and how fine they are.

(Transparency 9)

This is an interesting slide that I could not resist showing. Last year, I was working on thermoacoustics and realized from the thermoacoustic equations that a standard calculation for

varying grounds did not take all necessary factors into account correctly. This plot also serves as a typical example of a ground impedance curve.

The impedance gets higher, the lower the frequency you go. In this region, the impedance is inversely proportional to the square root of ω . If the soil becomes denser with depth then the imaginary part of the impedance, the reactance, becomes larger. In this region the reactance is inversely proportional to the frequency. This plot gives you an idea about how ground impedance depends on frequency.

(Transparency 10)

Reflection -- and lots of ugly equations. Again considering invariance along the surface, the k_R component of the plane wave along the surface must be the same for waves going down as the waves going up. The wavenumber times the sine of the angle of incidence is the wavenumber times the sine of the reflected angle. Numerically the incident wave and reflected wave have the same ω over c , so this gives us our law of reflection.

The law of reflection is substituted into the equations.

(Transparency 11)

Then the ratio of pressure to velocity is our impedance.

(Transparency 12)

Some special cases are listed on this slide. A rigid surface does not allow any vertical velocity. What is the impedance? The impedance is p/v ; if v is 0, the impedance is infinite. What happens to the reflection coefficient if the impedance is infinite? The 1's are negligible; the cosine θ 's cancel; and the reflection coefficient is one.

A pressure release surface -- an example is going from the ocean to air. No force is exerted by the air, so there is no pressure. If pressure is zero, what is our impedance? Zero. What is my reflection factor? Minus 1. The wave is turned upside down by reflection. This is an important special case.

One more special case is listed. What if now my angle of incidence is very, very small, grazing incidence? What is θ ? θ was measured from the normal in this case. This is the normal optical case. What is θ ? Ninety. What is the cosine of 90? Zero.

So the grazing incidence, then -- we find that reflection factor is negative 1.

We want to identify some problems for later. This is just a summary of that.

(Transparency 13)

Here is a person listening to sound from some source. There are two paths for the sound to get to him. If the reflection coefficient is negative 1, what does he hear? We have 1 plus negative 1. That gives you zero, assuming we are low enough so that these path lengths are the same. So grazing incidence over any finite impedance surface is going to give us zero sound. This is not observed.

Izzy Rudnick solved this problem. In fact, he adapted an analysis used for AM radio to sound propagation.

I point out here the analogous optics case: Lloyd's mirror. A point source, close to a dielectric surface produces a dark band next to the surface. This is the Lloyd's mirror effect.

Another related problem in the literature is the treatment of sonic booms and ground impedance. Onyeonwu hypothesized that grazing incidence is the same as a pressure release surface, and used pressure release boundary conditions in a wave calculation. We will see that his hypothesis is incorrect.

(Transparency 14)

We also have seen that the sound speed in that wave equation is not necessarily uniform. If you have a change in temperature with height then the sound speed changes. The higher the temperature, the faster the sound speed. There is also an effect on relative humidity that is relatively minor. You plug all this stuff in having to do with the molecular weight.

The molecular weight of 29 is close to nitrogen, and is that of water. Water has 6 degrees of freedom -- one extra degree of freedom. Nitrogen and oxygen have 5 degrees of freedom. This analysis then gives a linear temperature correction and an absolute humidity correction to the sound speed as indicated on the transparency. That absolute humidity correction is relatively small, so we principally are looking at temperature effects. The sound speed is not uniform, and the wave equation must be modified.

(Transparency 15)

Before we modify the wave equation, here is a graduate course in meteorology, everything you needed to know. In the atmosphere pressure varies exponentially with height, and the scale factor is between 7 and 9 kilometers. If you take a packet of gas and lift it up, is that adiabatic or isothermal? Probably adiabatic. If a sound wave with gradients over this small region is adiabatic, then lifting some mass of gas would probably be adiabatic also.

So temperature changes. What if its temperature changed in such a way that it was lighter than the rest of the gas around it? It is a hot-air balloon. What is it going to do? It is going to go up. So that says that, in nature, if I a parcel of gas adiabatically, then the temperature profile with height has to match the adiabatic variation. Otherwise, the atmosphere is unstable.

On the average, the temperature of the atmosphere follows the adiabatic lapse. Very high in the atmosphere, the sun's heating dominates and the atmosphere gets warmer with height. Near the earth surface though, the adiabatic lapse case dominates and guarantees that the sound speed is almost always higher at the ground than it is up above.

Then there is the day-night variation in weather conditions. When the sun shines on the ground -- which never happens here -- the ground heats up hotter than the adiabatic rate, and you get a lapse rate faster than adiabatic; it is hot near the ground, then cooler above. At night, if you

have clear skies -- which we know we do not have here -- radiation cooling occurs. The earth is exchanging radiation with outer space at 3 K. The temperature drops. This is why you can have frost on nights that are warmer than freezing. The air next to the ground cools, and the temperature structure looks like this. This is called temperature inversion.

The wind profiles are affected by viscosity -- molecular viscosity, but turbulent viscosity, due to turbulent flow over the ground. The profiles are also affected by the fact that there is density variation with height.

The equations for a medium with a density variation with height and with viscosity predicts that the wind profiles will vary logarithmically or close to logarithmically, depending on stability conditions.

The problem that makes atmospheric prediction difficult is that inversions are almost complete decoupling. Wind directions may vary by 90 to 180 degrees across an inversion.

(Transparency 16)

For the purposes of this talk we assume that the effect of a wind on near horizontal propagation is the same as if we took the component of the wind in the direction of propagation. That is, the wind at a height z times the cosine of the angle between the direction the wind is going in and the direction of propagation. This is an approximation that looks terrible, but is really quite good.

So now we have an appreciation for how the temperature and sound speed may vary with height. We need a way of calculating or indicating what effects those temperature and wind variations have on noise levels at a distance from the source. Perhaps the best way to understand this is by use of ray tracing.

I am going to be very careful to explain the limits of ray tracing. Recall that if you choose not to come back to the advanced session, you can be a consultant and make hundreds of thousands of dollars, providing answers that are indicative of what is going on, but which may be totally wrong.

(Laughter)

We are not arrogant, because our best calculations are only within a factor of 5. Our error bars are not 10^{-5} . But we do get annoyed when published results based on simple models are off by factors of 50 or 60. This is typical of the sorts of errors ray tracing can lead to.

(Transparency 17)

Start with a wave equation and allow the speed of sound to vary. That is, c is now a function of position. We want to find a solution. The solution we try is given here. There is an amplitude and a phase term. I know this is not a phase since it has units of distance, but it is called a phase anyway.

(Transparency 18)

The trial solution is substituted in the wave equation. Setting the phase equal to a constant gives surfaces of constant phase. That is a wave front. The direction in which this wave front is going is the normal to the surface. Normals are calculated by taking gradients. A small patch of phase front is illustrated. The gradient gives us the normal vector. It is not a unit vector, but it gives us a vector in the direction of propagation. That is unitless now.

Some of you who have had some quantum mechanics might recognize that this looks like the WKB approximation that is used for calculating transmission through barriers. In that case, the phase is written as an integral over wavenumber times distance. The WKB approximation retains one more term than the ray equations. There is your quantum mechanics for the day.

We assume that the phase and amplitude are slowly varying and that our frequency is high. It is sometimes said that ray equations are zero-wavelength equations. WKB approximations, by the way, have been used for predicting sound propagation. Recently there have been two papers by men who have the same last name, Li. They are both very good papers.

For this problem the two terms of order omega squared are retained. Three terms are neglected and lead to the relationships listed on the transparency. Look at the listed relationships. Can anyone here, with a straight face, tell me that they can, *a priori*, tell me when a problem is going to satisfy those conditions? It turns out that *post priori* you are not too certain either. I hope, in addition to giving you some physical rulers, to give you some mental rulers for measuring when the relationships are satisfied

(Transparency 19)

We rearrange the two terms we have retained and cancel some things out.

DR. MIGLIORI: The rule for a WKB approximation is that if you can ride a bicycle over the potential, it will work. As soon as you --

DR. RASPET: What size bicycle? I am going to supply you with the size of the wheel of the bicycle. We will have a drawing of the bike later.

The Eikonal equation results: the gradient squared equals the index of refraction. The index of refraction is analogous to the optical definition, i.e., the speed of sound is divided into a reference speed of sound. In acoustics, we do not have a speed of sound in a vacuum, so the reference c_0 is arbitrary.

Usually the speed of sound either at the surface or at the source height is chosen. If you make a few approximations you can show the validity conditions are that the fractional change in index of refraction per wavelength has to be less than 1. The fractional change in amplitude in a wavelength has to be less than 1. These are a little bit nicer than the conditions listed in the transparency, but further approximations have to be used to get them.

(Transparency 20)

To get a Snell's law out of this in 3-D the Eikonal equation is expressed in terms of direction cosines, and the direction cosines are expressed in terms of a path length, ds .

The magnitude of the gradient of the phase is index of refraction. Dividing the Eikonal equation by n produces an equation that looks like a direction cosine relation, $\alpha^2 + \beta^2 + \gamma^2 = 1$, where α is $1/n$ the component of the gradient in the x direction, and β is the $1/n$ component of gradient in the y direction, and so forth. This expresses the gradients in a direction-cosine form.

(Transparency 21)

The vector ds is an arbitrary direction can be divided into its components dx , dy , dz , and the direction cosines are just the derivatives of how much x , y , z change along the path.

(Transparency 22)

So the two expressions are equal and the given equation results. Just a little trickery, here: take another derivative with respect to s of all these equations. This is the definition of a directional derivative.

Stick that back in and, finally, we have the 3-D Snell's laws. Several research groups have written 3-D codes for sound propagation with c an arbitrary function of x , y , z . A ray is launched and followed through space. These codes are useful for highly strange atmospheres. Such analysis has been done by meteorologists.

(Transparency 23)

The atmosphere is pretty much layered, so most atmospheres can be modeled with variations in height only. (I could be fancy and say something about the horizontal scales being much larger than vertical, but I will not.) Then, on that previous page, all those derivatives with respect to x and y are zero, and the index of refraction times dx/ds , which is just our direction cosine in the x direction, is α_0 . The direction cosine in the y direction is constant. That means that if a ray is launched in a plane, it will remain in that plane.

Now the problem is two dimensional, and that simplifies calculations. Choose the x, z plane. The direction cosine times n is conserved. This is my direction cosine; α_0 times n is conserved. If a ray is launched at angle θ_0 , where the sound speed is $c(0)$, then the ratio p is conserved.

As long as the height variation is in the z direction, this is always a conserved quantity. The direction cosine in the y direction is the sine of θ_0 .

(Transparency 24)

One more heavy slide. If the form dz/ds equals the sine of θ_0 is put into the third Snell's law, an equation for the curvature results. The curvature is $\cos \theta_0$ over the speed of sound.

Here is a fudge. I worked down to here and I was getting the wrong sine, so I put an absolute value here. The absolute value of the curvature is the value p , c over the cosine θ_0

times the gradient at that particular height. The radius of curvature is 1 over that, and that is $1/p \times g$. So that tells me the local curvature, depending on what the speed of sound is.

There are some wonderful papers from the 1950s in which people have explored: Here is a form of $c(z)$; what sort of ray path do I get? There are three papers or more in which people calculated possible ray paths for different variations of speed of sound with height.

(Transparency 25)

The sound speed variations in the atmosphere are relatively slow. In a lot of cases, the atmosphere can be modeled as linear gradients patched together. An arbitrary curve is replaced by linear gradients and is constant everywhere -- it is a global constant -- g is constant in the layer; the radius of curvature is fixed, and the ray path is an arc of a circle. Arcs of circles are joined together to trace the path of a ray in the atmosphere.

As an example, data is used to calculate a radius of curvature. A 1 m/s change in 1 meter is huge. This might occur close to the ground, with winds. It would require a 1.5° C change in a meter to get this radius of curvature. The radius of curvature for this case is 340-meter . More typical would be 1000- to $10,000\text{-meter}$ radii of curvature.

(Transparency 26)

A ray launched upward will turn and go back down if the speed of sound at altitude is higher than $1/p$. At this point the cosine of θ is one and the ray turns downward. For any launch angle then, we only need to follow a ray if the speed of sound is higher than $1/p$ at some point; otherwise the ray can return to the ground.

(Transparency 27)

Now let us look at how we can follow rays through layers. You have two possibilities. A ray turns in a layer and comes back down. There is an analytic form for horizontal distance traveled. Analytic forms are given for how high it goes. In a computer problem, the only input needed is where it starts and what its angle is, and then where it leaves the layer can be calculated.

(Transparency 28)

If it does not turn, the same sort of calculations can be developed. Given the thickness of the layer, how far the ray travels through the layer can be calculated.

Another quantity that can be calculated is the time of flight. That is also a closed form. Another important parameter that can also be calculated is the range derivative.

(Transparency 29)

The ground is probably most easily modeled using a mirror atmosphere. That way you do not have to write an "if" statement; you just have infinite layers going in both directions. Some programs the reflection coefficient off the ground into account.

(Transparency 30)

Examine two rays; no energy passes across the rays. So the intensity at the receiver is the energy that is between the rays divided by the perpendicular area if atmospheric attenuation is taken into account.

DR. MIGLIORI: You mean, if you ignore attenuation.

DR. RASPET: Yes. We will ignore attenuation now and adjust for it later.

If the sound source was here and there was no refraction, the sound rays would travel in straight lines, out to this distance, and this would be the area. Due to refraction, there is a different area, and those two areas can be calculated. We would like to compare the spherical spreading. The comparison will be that the intensity ratios depend on the area, the intensity when the rays are refracted.

I have drawn the figure this way to remind you that the atmosphere is cylindrically symmetrical. This means that the area becomes smaller as rays are launched closer to vertical; a cosine theta factor. The areas are actually in a donut-shape. The area for straight rays is the radius, the horizontal radius, times 2π times $R d\theta$. The area for the refracted rays -- this is a normal area -- is $2\pi R \sin \theta_{\text{final}} dR/d\theta$. The range changes as the launch angle is changed. This trajectory is a bullet from a gun.

Changing the angle, changes the range. This is the range derivative. Take the component -- that is the sine of theta -- and it gives the perpendicular area, and then multiply by $2\pi R$. This gives the focus factor.

(Transparency 31)

The focus factor is just the perpendicular area without refraction over the area with refraction. If the rays are straight, the focus factor is 1. If there is a single layer with constant gradient, the focus factor is also 1. If there are multi-layers, the focus factor can differ from 1. What happens if the change in range with respect to angle is zero? What is the value of the focus factor? Infinity. I believe that is a large change in amplitude within a wavelength.

There are obviously problems with this $dR/d\theta$. What about a sprinkler? Where does all the water end up? All the water ends up at the end, where your time rate change of radius goes to zero; dR/dt goes to zero. The differential cross-section for that is infinite.

This is also the same as a rainbow scattering, which occurs where you get a reversal in range derivative. As the angle of incidence gets larger, the angle of refraction gets larger for a while, then comes back and gets smaller. It also occurs in nuclear physics, when you can run directly through a particle. There is a region where there is maximum deflection. At maximum deflection, the scattering cross-section is infinite.

Don't worry too much about the infinities, because you always have to multiply by a $d\theta$, and these things times $d\theta$ are finite.

DR. MIGLIORI: You get wet everywhere.

DR. RASPET: Right, and the water does not pile up on the edges.

That is one approach to ray tracing. I want to mention Fermat's principle.

(Transparency 32)

Fermat's principle is much more powerful, but you cannot get the validity criteria that are needed to know whether ray tracing can be used. Fermat's principle says that rays travel in a path which is stationary with respect to travel time. That means it travels in paths that are minimum paths of time, maximum paths of time, and inflection points.

This is neat, because you can derive the law of reflection from this. You also end up with some ray paths that we have not identified. Here is a circular object. Fermat's principle predicts ray paths that look like this. Fermat's principle leads to the geometrical theory of diffraction, which is an extension of ray tracing to include diffractive effects. It has been widely used in electromagnetics, and also in optics, and I believe Dr. Marston has used this for scattering off objects.

DR. MIGLIORI: Why is an inflection point included in this?

DR. RASPET: Because that is stationary.

DR. MIGLIORI: Not with respect to time, it isn't. It is only the minima and the maxima, the least action.

DR. RASPET: Yes, I believe you are correct. Let me think about that.

(Transparency 33)

It is now time for your laboratory. You can work as lab partners and split it up, or the smart one can do it and the other one can take credit. This is a simplified ray tracing, in that we are going to assume that the radius of curvature is constant. It does not depend on the angle you start with.

You have the instructions. Draw rays coming from the source using the following rules: Ray paths are circles tangent to the initial direction. The radii are given on the work sheet (this is counter to all your calculus books). Plus is downward curving, minus is upward. Angle of reflection equals angle of incidence at the ground. At a layer interface, the different circle radii are tangent to the same line. When I specify $R = \text{infinity}$, that is a straight line.

If you find some place interesting, you are not restricted to the ray starting paths I have indicated every 15 degrees. You can add additional rays as needed to understand the ray field. Here is the scary part: Think about what a listener would hear at the different positions.

(The participants worked on the stated problem.)

DR. RASPET: An interesting thing happened here. I put this together and had constant radius of curvature to make it simpler. Obviously, the radius of curvature depends on the gradient and the cosine of the initial angle. It turns out that Dr. Stern has a construction that is

just beautiful. It would never have occurred to me. You can do this right, and even have it simpler than the drawings we are doing.

Hopefully, Dr. Stern, maybe in a discussion group, will show us.

(Transparency 1)

Now we're ready to start looking at the advanced level. Once you hear this, you are ruined for life; you cannot be a consultant. We are going to look at problems with the above that we come into trying to use these constructions in plane wave reflection. Then we are going to solve that problem about real reflection from the ground, look at a general solution. It turns out that all these are related solution techniques. We will talk about them.

(Transparency 34)

Let us look at this picture. That is my version. There are actually two things on here. One I am going to point out now. If I drew in more rays down here, I would end up with a big concentration of rays close to the ground. You would expect the level to get very loud close to the ground. That is the first thing we observe here. I am going to put this one aside, because there is another thing hidden in here, too.

(Transparency 35)

Case II: What do you hear? Nothing. Yet, when you go outdoors, generally the ground is heated, we have upward refraction -- this is the normal case for human beings, crickets, whatever -- and yet we still hear something. So here is a case where ray tracing gives us zero for an answer, and we suspect that zero is not the right answer.

(Transparency 36)

Here we get more subtle. The first thing is, there is a bunch of rays hitting here. This one comes out there, lands here. This one comes up here, lands here. We have a case now where I increase the angle, I get more and more range, and then, all of a sudden, as I increase the angle, I am coming back. So what is $dR/d\theta$? It is increasing and decreasing. I have to get to zero somewhere. That says that somewhere in here my focus factor heads toward infinity. How many infinite sounds have you heard recently?

There are fixes for ray tracing. If you go back and do WKB, take another term so you calculate the amplitudes. These are called caustics. In the high-frequency limit, you get a $\pi/2$ phase change. That takes a pulse wave and makes it look funny. In other words, it is a phase shift of i . So here is another problem where rays give a wrong answer.

(Transparency 37)

What is the answer out here? Is this the same as case I? I got that every one of those escaped. That looks like a shadow zone. The graduate student working on elephant communications gave a desperate call. He was preparing a paper, and ran across this problem

and called me: "I ran this case that I know gives me zero for an answer, and I got a large value."

The problem here is that, if we look at the profile, the sound speed never gets higher than at the source, and so you do not expect to see rays coming back. But then you look at the arrangement, and if you know anything about ocean propagation, see that ducted propagation occurs along this axis. Energy is diffracted into the waveguide. The sound does not interact with the ground and there is no ground absorption. All that affects the sound is spherical spreading. Even the leaky waves, down here, are loud. This is a situation where elephants can find each other, even though ray tracing says it cannot be done.

Low frequency waves are sort of hard to model as rays because they have large wavelengths.

DR. GARRETT: I think a point of interest here is that that is not spherical spreading, but cylindrical.

DR. RASPET: It is cylindrical, yes. It is spreading out in that plane, and it is $1/r$ intensity, 1 over square root of r pressure. Even though the only way you hear it is leakiness and the only way you feed it is leakiness, when you get to large ranges, it is still quite loud.

So we have three "not-observeds." We talked about Lloyd's mirror. I mentioned that at those caustics, or when rays become tangent and we see this range derivative, we get a $\pi/2$ phase change. The two cases, I and III, violate our idea that the amplitude change in wavelength is small.

(Transparency 34 again)

I put aside the first one. I wanted to show you something from it, too. If we look at this mirror atmosphere here -- here is the ground -- then this looks like one of those cases where we have a higher gradient above a lower gradient. Caustics are also formed in this case. Caustics are formed for a single gradient above a reflecting surface.

There is a larger gradient over a smaller gradient; $dR/d\theta$ changes direction. When you write a computer program and search through layers, you find different behaviors and you can't do a linear search. Instead, you must find out where $dR/d\theta$ is zero -- that is the ray that goes like this -- and search below it and search above it, for two different families of rays.

(Transparency 36 again)

(Transparency 38)

Here is a summary slide.

(Transparency 39)

If you want to write a computer program, there are a few things to note here. The rays come back only if our speed of sound above is greater than our speed of sound at the source. We have to treat differently layers in which you have an increase in gradient -- and that could be

going from a negative gradient to a positive gradient, or from zero to a gradient -- since caustic formation occurs.

If a numerical answer is required, a WKB-type correction must be performed. We saw that we could treat the ground as a mirror. Dr. Atchley points out that this is not possible if you have an ocean bottom here, because then you have two mirror planes and you end up with an infinite construction. Luckily, with the ground and outdoor propagation we have a mirror and the calculation is usually terminated by the adiabatic lapse rate.

In writing a computer program, it turns out that, in addition to finding time of flight along the path iteratively -- in other words, you just launch and go through the layers; you do not have to have any interdependence of layers -- you can also calculate the range derivative iteratively and calculate the focus factor. To find out how you get from a source to a receiver, you use Newton's method for each class of eigenrays.

Then you get involved with fixes. Twice I have been involved in developing ray-tracing programs. In each case, we kept adding fixes until we had a computer program that we could not run. Ray tracing is a wonderful tool for looking at things, but is very poor in the atmosphere for predicting things quantitatively. At the long ranges that I am interested in the frequencies must be low. That means the wavelength is large, and fractional changes in index of refraction and fractional changes in amplitude become important.

We also see that, where rays give you zero for an answer, another calculation is needed. We do not observe zero levels; therefore, a better way of calculating levels is needed.

Back to equations. If we have cylindrical symmetry -- that means that the atmosphere is perfectly still and there is a variation in sound speed with height only or, to an approximation, there is an anisotropic wind but I assume that its effect is cylindrically symmetrical -- that is a big assumption.

Hopefully, each of you will investigate this carefully rather than taking my word for it, and publish a paper that solves this problem analytically.

(Transparency 40)

In this case, the wave equation and Helmholtz equation in cylindrical coordinates and a fixed frequency is assumed. There is a source at position r' and height z . Separation of variables is achieved with an integral transform. This is a Bessel transform involving the Bessel function J_0 and a transform variable K . I will explain what K is. Later, I plug this form for p into the Helmholtz equation.

(Transparency 41)

That is the solution for the real p as a function of r and z . Here is the transformed \hat{p} , and this is a differential equation for \hat{p} . This is a 1-D equation. In the speed of sound, this $K(z)$ is ω over $c(z)$, and this variable K . This K can be interpreted as the horizontal

component of the wavenumber. Notice I did not put "horizontal" up there, because every other day I convince myself that it is the vertical, not the horizontal, but it is correct that it is the horizontal component of the wavenumber.

From ray tracing, K is proportional to $1/p$. This is a conserved quantity, meaning that K stays constant throughout the propagation. This is a solution to a 1-D equation meeting the boundary conditions.

(Transparency 42)

A complex impedance ground used to calculate plane wave reflection. The ratio of the pressure amplitude to the normal velocity is a complex number. These are complex amplitudes. It is our impedance that we can look up from one of those 400 or 500 papers on ground impedance. As I say, this is a field in which you can publish the same paper every 20 years, and many people have.

(Laughter)

Pressure is continuous at the source. This is a discontinuity in vertical velocity. A point source has vertical velocity going up and vertical velocity going down, and so there is a two-times-normal velocity source discontinuity. At infinity there is the radiation condition.

Hey, this is a standing wave tube, an inhomogeneous standing wave tube that we are solving. An inhomogeneous standing wave tube with a source in it.

(Transparency 43)

Here are applications. I will not say who adapted the analysis for AM propagation to the ground first. The mentor got a few things wrong. Many other people got to publish on this.

(Transparency 44)

In spherical wave reflection, there is a point source with direct and reflected distances, and there is reflection from the ground. That is the ray picture. We have a uniform atmosphere, which means that the solution to the wave equation is just e^{ikr} . There are no changes in sound speed with height. We ask, how does this differ from plane wave reflection? I do not have some infinite plane wave source here. This is a spherical source. The angle of incidence changes along the ground so there is an amplitude change. When an amplitude changes, diffraction occurs.

(Transparency 45)

I have presented a mixed analysis here, to try to make this understandable. By the way, better ways of doing this have been published within the last year. Instead of integrating in K space, we look at taking that transform.

We want to integrate in wavenumber space from minus infinity to infinity. Brekhovskikh integrates in terms of complex angles. The contour in K space becomes this integral in theta space. As long as we do not pass a pole, we can distort the contour and make things easier.

What I point out here is -- this is just our reflection coefficient and things that come out of doing this -- here is the phase of our integral. The distance r is big for propagation outdoors. The phase is wiggling wildly. The further out, the more it wiggles. Is this an integral you want to do? Nobody wants to do this integral. But it is a very important integral.

(Transparency 46)

We have the freedom to distort contours in the complex plane as long as you keep track of the poles. So the cosine of theta minus theta sub zero wiggles. If we choose a contour where theta is complex, then the integral will vary smoothly. As a matter of fact, by complex variable theory, you are guaranteed that you can find a place where it does not wiggle at all.

The e^{ikr} portion of the integral will be completely real and decay exponentially. The choice for the integrand is a 45-degree contour through theta sub zero, the angle of incidence, and the contour behaves as illustrated. Which contour do you want to do, the wiggly one or the smooth one?

This contour is the steepest descent or a stationary-phase contour. It is both of them: The phase does not change in the exponent, and it is also the steepest way to get down from the saddle. This means that the whole integral can be evaluated by some small region around the saddle point, and an series expansion can be used.

What has led to some confusion over the years is that under some conditions the contour crosses a pole. The denominator of the reflection coefficient has a zero. When the imaginary part of the impedance is greater than the real part of the impedance, then you also have to take into account the distortion of this contour by that pole. Small complications.

We have solved a standing wave tube, and we do the integral. It is not easy, but I promise you that you can follow Brekhovskikh's explanation. There are 25-plus references on evaluating this integral in my cabinet.

By the way, if you can get hold of the old book, *Waves and Layered Media*, it is better than the two-volume set. The two-volume set threw away a lot of good things to stick in some pretty bad things. Also, Felsen and Markovitz is an excellent reference.

DR. MARSTON: Which edition of the old book do you like?

DR. RASPET: The maroon one -- a wonderful book.

(Transparency 47)

Donato breaks this expression up into easily understood pieces. He was a nuclear physicist with the National Research Council of Canada, who dabbles in acoustics. NRC has supported continuing basic research in sound.

The wave is broken up into a direct wave, a perfectly reflected wave -- and this is also a term that is dependent on $1/r$ -- and a ground wave. Let us look at the ground wave. It is proportional to $1/r^2$ in pressure.

Then there is the surface wave term, which has fascinated people for years. It depends on the square root of d , meaning it decays cylindrically, and it is odd. What a mess. This term only exists when the real part of the impedance is less than the imaginary part.

What happens as we go to grazing incidence? The direct wave cancels the reflected wave. These two terms are identical as grazing incidence is approached. Cosine theta goes to zero, so this goes to -2. All terms cancel except for the ground wave which decays as $1/r^2$. In E&M this terms explains why you can hear AM radio from Chicago in Mississippi and you cannot hear Chicago FM stations in down-state Illinois. This is exactly the same diffraction term that allows long-wavelength AM radio to propagate over long distances.

The surface wave is very peculiar. In your handouts (but not in my lecture) is a solution to a nice plane-wave/surface-wave problem out of Pierce, problem 313, in which he shows that there are conditions for having a very weird plane surface wave over a complex impedance surface.

Its direction of propagation is downward. horizontally. Donato's expression is just the cylindrical wave version of that surface wave.

(Transparency 48)

In a word, when the plane wave rays cancel, then our diffracted terms become important. Really, on this reflected wave, it includes that ground wave, which is due to diffraction. That orange portion is the ground wave. It is due to diffraction, because we have changing amplitudes at the surface, and we also have a peculiar surface.

The surface-wave path is a stationary path from Fermat's principle. The surface wave is always slower than any of the other waves. In fact, to be a surface wave, it has to be sub-sonic.

(Transparency 49)

This slide shows some calculations; this is normal sound propagation over a surface. You can see that close to the source where we do not have grazing incidence, the level is almost constant. Far away from the source in the asymptotic region, the pressure is proportional to $1/r^2$. In that region the direct and reflected waves cancel and all that remains is the ground wave.

There are much nicer forms than those of Donato for calculation that we developed by Hayek, Attenborough, and Lawther, while Attenborough was at Penn State.

(Transparency 50)

A second case illustrated on the transparency is the case where the imaginary part of the impedance is much larger than the real part. In this case the surface wave is excited. Close to the source the level is higher, but the level is lower. There has been speculation that the surface wave could lead to large levels at long distances. That is not the case.

(Transparency 51)

Next, I'll talk about the fast-field program. It is a full-wave solution, using the same trick we started the advanced lecture with. Assume cylindrical symmetry, then do a Bessel transform. Now I want to solve up and down motion and then transform that to find out what the horizontal motion is. The FFP was developed for underwater acoustics. We put in complex impedance on the bottom.

(Transparency 52)

Just like the experimentalists here like to build their own equipment, we like to build our own computer programs so we know what is inside them. We did that and put complex impedance ground into the FFP.

So we have a point source in a layered atmosphere comprised of uniform layers. There are programs written that use gradient layers. The problem with this is that the solutions are in terms of the Airy function instead of sines, cosines, and exponents. Calculating uniformly convergent Airy functions can sometimes be unpleasant.

We solved this by taking propagator matrices, going up and down, like Dr. Maynard did in the string.

Another method -- I was working with some electrical engineers, and they had a very nice program for doing transmission lines that automatically terminated. One of the problems with propagator matrices, with this sort of atmosphere, is that you are going up and up and up and attenuating things, you meet a boundary condition and you start coming down, and you can amplify noise, amplify the computer noise. So it is nice to have tricks for terminating things.

Richards, at the University of Houston, had a very nice program for solving transmission lines, terminating them, and getting nicely convergent solutions. We based one version of our program on that and got nice results.

(Transparency 53)

Now we are back to evaluating the integral. This is a very general case. We do not have one stationary phase point. We may have more than one pole. We cannot use complex variables to evaluate this. Instead, we take this and we express J_0 as the sum of the Hankel functions of the first and second type.

Examine the asymptotic behavior of the second type, and it is a wave coming towards the source. With some sort of belief in causality, we throw that one away. There are actually stronger arguments that exist in the literature for throwing it away that have to do with using the method of steepest descent close to the source. There have been papers saying that sometimes you need to keep this, but these only apply to pathological cases.

Asymptotically, the integral is proportional to e^{iKr} . This is exactly what was done to get the solution for spherical waves. when this is substituted the integrand is in the form of a Fourier transform. Discrete FFTs can do this integral rapidly.

The Fourier transform gives the solution for several different range points, like 1024, 2048, 4096, in one operation. This is faster than doing the transform directly. The FFT goes as $n \log n$ instead of n^2 , so there is significant savings in time. The "fast" in fast-field program is sort of laughable. Depending on the problem, this may take two minutes to two hours. The lower the frequency, the better it works.

DR. MIGLIORI: What do you do about reflection coefficients?

DR. RASPET: This is an impedance boundary condition, so that is all included. The propagator matrices are either written as a traveling wave going up and down, with reflections at each layer, or as a standing wave in each layer. You can write it either way. Those are all taken into account.

DR. MIGLIORI: Are the layers always assumed to be bigger than the wavelength?

DR. RASPET: The layers are always assumed to be smaller than the wavelength. A very good question.

DR. MIGLIORI: When you switch over, you are in trouble, right?

DR. RASPET: Right. It appears to work for half-a-wavelength layers. You rarely get into trouble with that. Occasionally, with rapid changes, sometimes you need more. That is a very good question. This is why people looked into computing Airy functions, so you could model the atmosphere with much bigger layers. It is one of those things where you trade off one kind of complexity for another kind.

DR. MIGLIORI: One of the failures of the WKB is that it does not include any reflections from gradients in the potential.

DR. RASPET: Right.

DR. HAMILTON: You mentioned that the surface wave can be important. How does that affect the normal impedance assumptions that you use in these codes?

DR. RASPET: This surface wave is the complex impedance surface wave, which is due to the assumption of having a normal impedance surface with the imaginary part bigger than the real part.

The clean way to look at this surface wave is the comb-like surface. If you have a surface that looks like an upside down comb the sound can penetrate it along the normal direction. The travel path forces a phase delay. So the impedance has an imaginary part much larger than the real part. Beautiful surface waves can be set up and measured for such a surface.

The people at NRC have done this experiment. They take light fixtures with egg crate-type holes, put them down on the ground, and propagate surface waves over long distances and measure all their behavior. There is not an elastic surface wave; it is a surface wave over a complex impedance surface. Brekhovskikh has a nice paper on this, and Tolstoy has also measured it.

DR. MIGLIORI: Is this the equivalent of the near-field antenna problem, electromagnetic antenna problem, so in near-field you have things that sort of propagate, except they die out rapidly?

DR. RASPET: The problem I spoke about, whether you keep H_0^1 and H_0^2 , is a near-field problem. It is sort of an extension of what we are talking about with rays. As long as you have significant propagation of the horizontal direction, then you do not need the second solution. If you have something like a pressure-release surface so you have very little propagation horizontally, then you do need it because you have screwed with the steepest descent type of solution.

(Transparency 54)

Here is an example of a fast-field program. This is a plot of sound speed versus height. There are quite a few more layers than in the ray laboratory, so it would take a while to draw the rays for this. Look at the structure. This represents a low level temperature inversion. The gradient gets steeper, so we expect to see a caustic. What sort of solution would you expect to see?

If I have a 100-Hz signal and I am standing out 200 meters, 400 meters, would this be loud or quiet? The source is on the ground. There is a male elephant here, and a female elephant there. You would expect it to be loud. If you look at your ray pictures, you would expect it to be very loud in some places -- $dR/d\theta$ is small. What do we really see?

(Transparency 55)

First, we have to look at a solution to the vertical problem. That is our p -hat that we are taking the Fourier transform of. What does this look like? From someone who works in resonant ultrasound spectroscopy, what does this look like?

DR. MIGLIORI: You do not have any axes labeled there. So you have frequency and amplitude. It surely does not look Lorentzian to me.

DR. RASPET: It looks like a resonance curve to me. This looks like a resonance curve of a vertical standing wave. This is the resonance curve for the atmosphere. The resolution is poor. There are many modes in the single peak.

DR. MIGLIORI: I do not understand.

DR. RASPET: When we have a Fourier-Bessel transform, we solve the vertical wave problem for different wavenumbers. When the boundary conditions are such that I get resonance, I am going to get a large response --

DR. MIGLIORI: You have discrete little bumps in there, so you do get --

DR. RASPET: There are different modes in here.

DR. MIGLIORI: You do get resonances in the vertical direction. But if it were continuous, you might not.

DR. RASPET: I will always get one bump in this, because I will always have a mode that propagates out from source to receiver.

DR. MIGLIORI: You are going straight up now.

DR. RASPET: This is the straight-up propagation at wavenumber K . K is the horizontal axis and p -hat is the vertical axis.

DR. MIGLIORI: If instead of having discrete breaks in the gradient of the sound velocity you had a continuous pattern, you would only have one --

DR. RASPET: No. I would still have different modes. In general, there is one propagation speed corresponding to horizontal propagation. There are some launch angles that go up, come back down, go back up, that give us integral solutions to our top and bottom boundary conditions. Each of those will give a solution here. This is, of course, in the complex domain. We only display the magnitude.

Let us get to real things.

(Transparency 56)

Here is the solution. This is the sound pressure versus range. This is not excess attenuation. The ray picture would give us cylindrical plus some extra bumps. Generally, it would be cylindrical propagation, and then we would have some focuses wherever $dR/d\theta$ is small. Our focus regions, which on your ray picture were small, are now spread out. The sound pressure decays slower than $1/r^2$. Ground interactions eventually begin to cause absorption out here.

DR. MIGLIORI: If this were just free fields and nothing interesting was happening, it would be $1/r$, right?

DR. RASPET: This is intensity, so it would be $1/r^2$.

DR. MIGLIORI: You have pressure on the left, right?

DR. RASPET: Yes. Thank you. I was very bad. I knew that I was, but I could not catch them all.

DR. MIGLIORI: So $1/r^2$ is what you expect if nothing good happens.

DR. RASPET: If I am in a homogeneous atmosphere, far away from all boundaries, I expect to see $1/r^2$.

DR. MIGLIORI: But that should really be $1/r$ on this plot, because you are plotting sound pressure level. I am not trying to be difficult; I just want to be sure I understand it.

DR. RASPET: You are doing a great job of being difficult, of doing what you are supposed to do. I truly appreciate it.

DR. MIGLIORI: Why did you plot square root of r ?

DR. RASPET: I did not plot square root of r . Those lines are to indicate what the slope would be if pressure were decaying. It is 1 over the square root of r .

DR. MIGLIORI: So 1 over the square root of r would be if it were confined to a waveguide layer.

DR. RASPET: Right.

DR. MIGLIORI: That is the best you could ever hope to do, and $1/r$ is the worst you could do.

DR. RASPET: You could do a lot worse than $1/r$ if you have upward refraction.

DR. MIGLIORI: So those two bound the 2-D and the 3-D propagation problems.

DR. RASPET: Yes. This is a nice illustration, when you put numbers in and start looking at these longer wavelengths, the ray picture is a good indicator of what is happening -- the levels stay relatively high; we get some increase here -- but the ground affects things. The caustics or the focuses are not very small; they are spread out. So we need to be careful in taking ray pictures too literally.

(Transparency 57)

Here is an elephant call, 15 Hz, at three times of day, with some idealized profiles. David Larom took data over three months in Namibia on meteorological conditions -- this is from the paper that has idealized profiles -- and then ran every profile he got. He called me up out of the blue -- he is at the University of Virginia -- and said, "Can I use ray tracing to predict what elephants do at 15 Hz," and I said a few foul words. He came and visited and learned how to run the FFP and has done some very nice work.

Here are daytime effects. Daytime heating is the orange. Adiabatic lapse rate would be some sort of transition time. Then at night the radiation cooling occurs and an inversion forms. These are the decays. The plot displays attenuation not excess attenuation.

The estimated threshold of hearing for the elephant is 73 dB. An elephant produces 140 dB at 1 meter, so -67 dB down from 0 dB at one meter corresponds to the estimated threshold of hearing. In the daytime, an elephant call is audible only to 2 kilometers. This is not going to do him any good if he needs to communicate 6 to 10 kilometers.

In the evening, as the lapse rate lessens, the elephant communication range is only 3 kilometers. But in the evening after the inversion forms the elephants can communicate at 6 to 10 kilometers. Elephants have responded to calls at about 7 kilometers. We would like to do measurements, of course, to see whether these calculations are meaningful. In any case, they are nice calculations.

DR. MIGLIORI: What happens early in the morning when the inversion layer starts to move up?

DR. RASPET: You get a waveguide formed. That is the time of day when military bases get the most complaints. Early in the morning, as ground heating begins, an elevated waveguide

forms. That is actually when sound levels are loudest. This waveguide is very unstable and not very useful. The elephants do a lot better communicating in the evening.

(Transparency 58)

We made some measurements to compare to the fast-field program. This is 18 months of my life on one sheet of paper, in which we got up at random times of day every weekend for 18 months and took data. From the top, there is upward refracting -- those are your rays going up -- slightly upward refracting, neutral, downward refracting, very much downward refracting. These are comparisons of measurements to data.

You can see how unsuccessful even the fast field was. These are the fast-field program results. So for upward refraction, you get reasonable results in the near shadow, but then the predicted levels drop way below the measured levels.

So the FP, compared to our data worked very well for neutral and downward refractions, but was very poor for upward refractions.

I want to point out that our latest turbulence calculation is included on this graph and matches the data. But the interesting thing at this point is that the fast-field program has some problems.

(Transparency 59)

The questions that arose: What is this accuracy for upward refraction? Do we have a problem in our prediction or is this an additional physical mechanism? Another thing that came up was, people looked at that surface wave -- as I say, they have a physical identity there -- and they said, that surface wave is running along the ground; maybe it is not affected by refraction, maybe it is going to tunnel through and give us a solution. This was also an interesting question: What is the role of the surface wave in upward refracting?

Finally what is the criteria for enhanced downwind refraction? We saw that for downward refracting cases, the FFP did not predict the huge gains we expected. It is obvious that the cause is the interaction of the sound with the ground. We were curious as to whether we could develop an analytic calculation to compare with the FFP.

(Transparency 60)

Steve Franke at the University of Illinois at Champaign-Urbana developed a residue series solution for upward refraction and wrote convergent Airy function routines. In this case we pick a special atmosphere. This is a continuously varying atmosphere. For small radius of curvature, this is about linear. This is a way to approximate linearly upward refracting.

DR. CHAMBERS: Don't you mean when the radius of curvature is large?

DR. RASPET: Yes, when the radius of curvature is large. Thank you, Jim. The radius of curvature is large, so that the sound speed profile is approximately linear. Plugging this $c(z)$ into

the Helmholtz equation gives the solution for vertical propagation. So this is a solution for a special profile.

Again, this is just like a tube that has a change in sound speed with length along the tube. The solutions are Airy functions. From quantum mechanics, you may remember that Airy functions are used to join together exponentially decaying solutions and oscillating solutions, because they decay on one side of the original and oscillate on the other side.

For upward refraction, the sound level decays away from the ground. As we go away from the ground, you should expect to see levels get smaller. We use the decaying side of the Airy functions. Now the solution spans all the way to the ground from the top of the atmosphere. No layers are needed.

(Transparency 61)

Our cost is that we have a lot of ugly equations. This term that I have crossed out is a direct ray. Far enough away from the source, ray tracing says that the direct wave should not contribute anything, so I throw it away. Look at this term. With the idea that we are looking for a resonant solution, we search for a pole to get a large contribution.

If you search for where the denominator is zero and evaluate that the residue (regular complex variable analysis), then you get a wave contribution, a single contribution. Pierce discusses this for the simplest cases and does a very nice job.

I want to point out the length ℓ here. This is what Pierce calls a creeping-wave-layer thickness. Let us just call it diffraction thickness. It is a radius of curvature divided by the wavenumber squared to the $1/3$. Another parameter in here is this factor q , which has something to do with wavenumber, the layer thickness, and the size of the impedance.

If you have a rigid surface so the impedance is infinite, q would be zero, and you would just get the zero of the Airy function derivative. If q is infinite, then we take a ratio here and the zero of the Airy function provides the pole. Between those cases, q is finite.

Pierce suggests a way of doing the search, expanding from the rigid limit, expanding from the pressure-release limit. I started doing these searches and my computer got lost. First let's look at how a creeping wave behaves.

(Transparency 62)

The simplest solution that a lot of people use is the creeping wave. Onyeonwu, whom I mentioned early in the lecture, said the ground should behave like a pressure-release surface because this illustration resembles grazing incidence.

It turns out that near to the earth's surface, inside the distance l , this solution is not valid. The solution is also used over curved surfaces. Upward refraction over a flat surface is the same as straight diffraction over a curved surface.

The sound follows a rather peculiar path, much like the surface wave. Do not take the path seriously. The phase of the solution corresponds to a ray along this path.

(Transparency 63)

Let's look at the pole search mentioned above. I set up a program to do this pole search, expecting to see things go from zero of the Airy function, to the zero of the Airy function derivative by some detour. Look what happens. I kept running these things. I would start one here, start one here, and expect them to meet, and this thing would wander off in a complex plane. This is pretty neat. You cannot always take what you read to be the truth.

It turns out that the normal search is distorted by the presence of the ground reflection pole.

(Transparency 64)

To put it in terms of wavenumber, it looks like that. Exponential loss is determined by the imaginary part. In the residue series, we normally only have to take the first term. In some cases, the surface-wave pole can be a contribution, but normally it has much more loss than this first pole.

(Transparency 65)

I worked very hard to get a very unnatural surface to do this calculation, so take this with a grain of salt. This is what can happen. These are different modal contributions for upward refraction. The first mode dominates everywhere except very close to the receiver, and then we get a small region where that surface wave condition gives us a large amplitude. The source couples in very well to that surface wave, but then it attenuated rapidly.

The comparison of the residue series with the FFP led to some surprises. The surface wave contributes to the residue series. The pole search wanders away from our normal search routine suggested by Pierce. The surface wave is only important in the transition region.

Comparison also showed that the FFP predictions are good down to about -120 dB. This is our numerical noise floor in the fast-field program.

That still leaves us with a big gap between prediction and data. That big gap is turbulence. As I've said, recently we have managed to get a better handle on that. Ken Gilbert at Penn State has done some very nice calculations of scattering into the shadow zone and how that increases the levels in the shadow zone. But that is not the subject of the talk.

DR. MAYNARD: I am having trouble correlating what you just said with Figure 6, the residue series solution.

DR. RASPET: I had a note to myself to pull this one out. What we found was that, when we took the residue series and did it correctly, that is what the residue series looks like. This is an excess attenuation plot. That was -120 in actual attenuation. This -50 or so here, where the residue series and the prediction differ, is actually 120 dB in absolute level.

DR. MAYNARD: That is what I do not understand, the FFP points, those points up above.

DR. RASPET: The physical factor -- this is a refraction prediction with no turbulence. The data points there are this. This is the turbulent FFP calculation.

It turns out that when you are trying to do the calculation with the FFP, you are taking contributions from lots of different wavenumbers and adding them up. To get a very low level out of that, you have to get the phases exactly right. We are adding numbers on the order of 100 to get a number that is on the order of 10^{-6} . That is very difficult for a machine to do.

When the turbulence decorrelates the contributions from the different wavenumbers, the different wavenumbers have the correct magnitudes, but we are scrambling the phases to account for turbulence, so we can do that accurately, even though the FFP will not work for the coherent case. It was somewhat of a revelation. At least that is our theory.

(Transparency 66)

One more solution. Once you have Airy functions running, you look for other applications. Getting these complex things to run nicely is a problem. Another application is downward refraction. Now the sound speed increases with height. Again, for small r , it is linear, so you can compare it to things like ray tracing, and you have the three boundary conditions.

(Transparency 67)

The solution is, to the untutored eye, identical to that for upward refraction, but, instead of decaying with height, this is wiggling with height. This is the other side of the Airy function. These are very strongly bound normal modes. Now the sound reflects back and strong standing waves occur. Each of those standing waves will give a large contribution to the wavefield. Again, there is a wave-layer thickness, ℓ , that is our ruler for standing under when diffraction corrections are important.

(Transparency 68)

Again, a pole search is executed. Again, one of the terms wanders away from its normal position. This is, again, the surface-wave term. In terms of wavenumbers, this is what the wavenumbers look like.

(Transparency 69)

You start out here with poles that have very little loss. The loss is increased because every time the sound wave interacts with the ground losses occur. There is a surface wave under some conditions.

For high frequencies, an ungodly number of modes are needed. The calculation cannot be done at high frequencies. At 100 Hz, for moderate gradients, there may be 50 modes. For 1000 Hz, 20,000 or 30,000 modes may be needed. But now ray tracing will work, because the wave lengths are small.

So normal mode and ray tracing are complementary. You have to use some sort of wave solution if you are within that wave layer of the ground. That is our ruler. Rays are not going to work inside that wave-layer thickness.

(Transparency 70)

Here is a case I ran. This one is ridiculous since it has a very low radius of curvature but it still is a good illustration. The important thing in these results is the loss incurred by the lowest mode.

The first mode of propagation is confined to the lowest 14 meters from the ground. It has an attenuation that will give you a $1/e$ length of 2000 meters. The ground gives us a rather large absorption for the first mode. The interaction with the ground causes large attenuation of these low-lying modes.

What if the inversion is only 20 meters high? Only the first mode propagates and it has a large attenuation. This explains some of the data, where downward refraction occurs and yet no enhancement occurs. The first mode interacts with the ground, loses lots of energy into the ground. Even though you have an inversion and you expect large levels, you will not observe it at long range, because this mode is so strongly coupled to the ground.

The other interesting thing is that this complex impedance displaces the mode so that the next one is about 46 meters. That is a very high inversion. A hundred feet, 200 feet, is a normal inversion thickness.

(Transparency 71)

We compared the normal mode strength to the output of the FFP, run for the Airy function gradient. This is just a histogram. The neat thing you can learn about the fast-field program and its spectra is that, just like a resonance, the loss is proportional to the width, and the area under the curve gives you the strength of that particular contribution. It is nice to compare this analytic solution to the FFP.

STUDENT: When you discuss the wavenumber there, is that a spatial wave number?

DR. RASPET: Remember, at the source level, you calculate what little k is, ω over c , and then times the cosine of θ . So the wavenumber is the horizontal component of the wavenumber at the source height. That is conserved because it is proportional to $1/p$.

(Transparency 72)

An interesting graph here. Recently, I worked with some people who had written a very nice ray-tracing program and compared it to the FFP output for certain curves. You know that there is no ugly baby to a mother; my co-authors were very pleased with these results, how good the comparison is.

I sort of thought the comparison was not that good, but it is interesting. At long ranges, the low rays within the wave-layer thickness are absorbed, and then there is a good agreement between the two solutions.

Close in, the low-lying rays are significant contributions, and the two do not agree well. It is a good illustration of the wave-layer thickness as a measure of whether ray tracing should work. If the region of interest is within wave-layer thickness of the ground; then you have to be very careful about using rays.

(Transparency 73)

So we have compared all these: normal mode versus spherical wave versus FFP and have learned a few things. The impedance pole always contributes. It is not always a standing wave. Sometimes it is faster than the speed of sound at the ground. It is only a true surface wave when you have mild upward refraction. The normal strength and attenuation can be related to the FFP spectrum. This helps us understand what that FFP spectrum contributes.

Finally, low inversions do not lead to large enhancement. That is very important. If they did, then every morning would be much noisier. Many military installations would be closed, because the military makes a lot of noise early in the morning.

(Transparency 74)

I have already said these, but I will put them up so you will know it is the conclusion. The full-wave solution is needed at caustics, in shadow zones, and near to the ground. For near to the ground, the number you need as a measure of the validity of ray tracing is the wave-layer parameter l . The surface wave behavior is interesting, but does not lead to enhancements at long ranges. The two analytic cases we have studied have helped us to understand how the FFP works. It would not have been possible for us to do the calculation with turbulence, unless we had done the two calculations with upward refraction and downward refraction, so that we could really understand the FFP.

Now your job is to complete the test and see if any answers are different and turn them in to me. Any questions?

DR. KEOLIAN: A lot of what you have talked about is similar to what underwater acousticians talk about. You guys seem to share numerical codes, perhaps, and analytic stuff. Between the two communities, what things overlap and what things do not overlap? Can you use some of the Navy codes for underwater propagation in atmospheric acoustics?

DR. RASPET: Yes. There is considerable overlap, obviously.

DR. KEOLIAN: And what does not overlap, what is different about atmospheric acoustics?

DR. RASPET: I think probably the codes would all run for both cases. The thing we have to be careful about is the ground impedance. We live close to the ground, whereas submarines

wander all through the medium -- probably the submarine captain who is very close to the ground would not be a captain any more, if he got there accidentally. That is the big difference. There may be convergence problems with undersea codes when we get close to complex impedance grounds.

The other thing is that ocean acoustics can be dominated by the surface, which is pressure-release and very clean. Experimentally, of course, sound propagation in the ocean is much longer-range, and so it is more important that underwater codes handle huge ranges. Our meteorological variability is so great that coherence is shot at relatively moderate distances. The coherent calculations do not make much sense at long distances.

Refraction is refraction anywhere; it is just a matter of what sort of boundary conditions you have and what implications those have for your codes.

DR. GARDNER: I would think that another major difference would be the Doppler. You could have winds that are a substantial fraction of c , whereas in the ocean you would only have currents that are not a substantial fraction of c .

DR. RASPET: Yes, I agree with you. I am willing to discuss this in a discussion session, but it is a 10-minute answer, and it is a pretty convoluted 10-minute answer.

DR. MARSTON: I missed some of the discussion. You indicated that with this downward refraction situation, the residue series was rather ineffective because you could have, presumably, many, many poles.

There is an analogous problem in scattering, where, if you were to calculate all the residues, it would just be too many, and it is not a good approach. But there are ways around it, some of them based on what is usually referred to as the Poisson summation formula. This kind of approach gives something still better than a simple ray picture. Is there any intermediate approach that has been worked out yet?

DR. RASPET: The things I have seen for sort of combined normal mode -- I do not know how you get the summation, but there are some things you can do with complex variables with distortion. Instead of summing up the poles, you look at distorting contour to make the summation shorter. I have seen people do that, but I have not.

DR. MAYNARD: Has anybody tried finite element codes?

DR. RASPET: Vic Sparrow is doing, not finite element, but finite difference, numerical, with some of them. I have not seen anybody do finite element. It is probably hard to match the boundary conditions. I do not know.

STUDENT: You did not discuss any range dependence or any range variation. Is that because --

DR. RASPET: These all assume cylindrical symmetry, so they are layered atmospheres. It works pretty well. The 3-D FFP listed in the graduate section of my outline does take into

account anisotropies and non-range-dependence. Range-dependent problems are usually worked using parabolic equation techniques.

STUDENT: You do not consider those needed for those ranges you are working on?

DR. RASPET: There are other, more competent people who understand them better, who are doing the research, such as Ken Gilbert. There are not enough people in this field for us to fight about domains. I wish I understood more. There are some nice range-dependent things where you can --

STUDENT: I was just thinking that, for the elephant studies, you were not looking at ranges that were all that variant to be concerned with that.

DR. RASPET: Meteorologically, no. The ground, obviously, when you start getting roughness and that sort of thing, has high range dependence. Mike White, at the Construction Engineering Research Lab, has done a lot with range dependence. But the FFP is hard to stop and restart.

SENSOR PHYSICS: SIGNALS AND NOISE

THOMAS GABRIELSON
PENNSYLVANIA STATE UNIVERSITY

DR. GABRIELSON: The Physical Acoustics Summer School is a new experience for me. I figured what I really needed was some relaxation, so I thought I would accept the invitation to talk about physics to an audience of people who understand physics better than I do. That has been a really relaxing thought in my mind lately.

(Laughter)

It is not the only reason I would rather be in the audience today. I could, frankly, use the sleep. But we will carry on.

(Transparency: PASS 96 Contents)

Professor Bass said that there was no logic in the schedule. That is certainly true if they invited me, in the first place. I am going to try to explain some logic in the topics that I pick, although that is going to be pretty difficult, too.

Sensor physics is really a fascinating field to me personally, because it integrates a lot of concepts in basic and applied physics and electronics; a lot of different things come together when you study sensors and how they behave.

I wanted to talk about some things that are not emphasized much in the usual books and classes. There are a number of books on sensor and transducers and many of those either treat some of these topics in a cavalier manner or ignore them altogether. But I find the topics interesting, and that is enough of a reason, since I got to pick the topics.

The first part of this talk is going to focus on self-noise in transducers. I am going to talk about two flavors of self-noise in transducers (equilibrium and non-equilibrium noise). The thing that sets a limit in using a sensor is frequently its self-noise, especially in the business that I come from. I am going to talk about the varieties of self-noise and from the start of equilibrium-thermal noise to the end of non-equilibrium noise, you will notice something.

Starting with equilibrium-thermal noise, I am going to talk about global principles. These are things that shape physics as a whole. We are going to discuss the second law of thermodynamics. Also, the principles tend to be rather simple to apply to a transducer. As we get down farther and farther in the discussion, the principles will become less global, the restrictive assumptions will build up, and it will be harder to apply some of these things, until we get to the burnout point, and then I will switch topics, because I cannot handle it any more.

Some people would argue that I could keep going on noise for the entire session, but I did not want to. I gave a noise talk at Penn State at the ONR Materials Conference and that night I

went to a restaurant in town to try to get away from the crowds. I walked past a table and I heard, "There's the noise guy."

(Laughter)

I am truly fascinated by noise. There is a lot of attraction to it, but I do recognize that there are occasionally, in rare instances, times when signal is important.

(Laughter)

So I am also going to talk a little bit about signal transduction in terms of sensor calibration.

(Transparency: Quiz Question 1)

My quiz is not like Dr. Raspet's. I am not going to give the quiz twice and collect answers before and after, because I do not want to be humiliated. But I thought I would start off with something to provoke a little bit of thought.

Consider a simple mass-spring oscillator. This is a mass hanging on a spring. It has a resonance frequency of 100 Hz. It has a Q of 10. It has a mass of 1 mg. If you displace that little ball by 1 mm, how long does it take for the initial vibration to decay in amplitude by a factor of 10^8 ?

A factor of 10^8 in amplitude takes you down to a tenth of an angstrom. I included some sheets following this slide on Q that I am not going to talk about -- they are just for your reference. Several people have already talked about Q: it is an important thing to understand, and you ought to do this calculation. Find out how long it takes to decay to a tenth of an angstrom.

Some people complain that this level is too small, but it is not too small at all. If you were to make an accelerometer based on this kind of a system -- a moving mass -- and measure the displacement of the mass with respect to the case, such a displacement is equivalent to an acceleration of half a micro-G. In my business that is a large acceleration. So this is not an outrageous number to consider.

So now that you have finished the calculation, I will tell you the answer. It never gets to that level. It does not get down to this level no matter how long you wait, as long as you are operating the system at room temperature. That is because there is molecular agitation forcing the mass to move. You can get it down to 0.1 \AA if you cool it to a few kelvin. That is not normally a practical way to make an accelerometer if you are using it for structural measurements or something like that.

(Transparency: Proper Dynamics)

This is a really important point about transducers. I would guess almost none of you have been taught the simple damped mass-spring oscillator correctly. It is kind of a gripe I have,

because this is a good opportunity to integrate some very important concepts in thermodynamics and dynamics, and it almost never done in education.

This system (on the left) has a mass, a spring, and a damper. In dynamics you glibly write the equation for the dynamical response: you hit the mass and it decays like this (a decaying sinusoid). If you let it go long enough, it will get down to any small amplitude you want to specify. It just keeps decaying and decaying. Nonsense. It does not do that.

The mass is made up of molecules. The spring is made up of molecules. It has molecules hitting it from all sides and the collisions force the mass to move. If the amplitude really decayed down to virtually nothing, that would imply that the molecular motion in the structure was decaying down to nothing, and this would be one heck of a better refrigerator than thermoacoustics, because it would cool anything really fast.

(Laughter)

Well, it does not work that way. This equation (on the left) violates the second law of thermodynamics in a real physical system. It is a mathematical equation; it is fine in mathematics. You cannot apply it to a physical situation.

The equation the right is the proper way to write the dynamics for this oscillator. It has a forcing function (F_n) and this is related to its connection to the environment. This forcing function, which is related to the molecular agitation, is a function of the damping, R (we will see how).

If you reduce R to zero, then you can reduce the forcing term to zero, because you have decoupled it from the environment. It does not mean that the oscillations will go to zero; it just means that they will not decay any more. But if there is an R in the equation, then there is a fluctuating force in the equation -- this is required by the second law of thermodynamics -- and this is how the dynamics equation should be written.

The mass starts off oscillating and then it decays into a noisy background. That is very important for sensors. It is not very important for the springs on your automobile. I can calculate the noise floor for your automobile on its springs, but it is not of much value. If I can figure out how to make it valuable, then I will sell it to the automobile companies and get rich, but noise background is not an issue in automobile suspensions.

(Transparency: Fluctuation-Dissipation Theorem)

There is a global theorem, called the fluctuation-dissipation theorem. I think it is kind of cool because the words are really big and you can also get at the meat of it just by talking rather than writing equations down.

Picture a system like this: this could be a microphone. This would be the diaphragm of the microphone and it would be hung on these little leaf springs so that the plate can move back and

forth. Maybe it would be a capacitive microphone with one electrode on the moving plate and one electrode on the fixed plate.

Take a look at what happens in the structure. As the plate moves back and forth, it squeezes fluid out between the fixed plate and the moving plate and there is some viscosity associated with that fluid, so there is some damping. It turns out that is a pretty substantial component of the damping in a small moving-plate system like this.

But there are other losses, too. There is mechanical damping in the structure where the spring and the plate connect, in the spring, in the connection to the wall, lossy bending in here, lossy bending in there. Then there are other things, also. The plate radiates thermally -- not very important at room temperature but could be important at elevated temperatures.

Also, when the plate moves, it radiates acoustically. There are lots of ways for energy to leave this system. One important thing about the fluctuation-dissipation theorem is to realize that the name is slightly misleading. You have to take the point of view of the system that you are considering. If you pretend to be the membrane of the microphone, anything that causes energy to leave you, you see as a dissipation.

You might be a loudspeaker and your goal in life is to generate acoustic radiation, but the radiated energy is LEAVING the system. In this sense, it represents a dissipation. That is the essence of the fluctuation-dissipation theorem. The dissipation refers to energy leaving the system, even if it is not what you would classically call dissipation. Not too many people call acoustic radiation dissipation, but from the point of view of the plate, it is.

The fluctuation-dissipation theorem essentially says that any time there is dissipation in a system, there are also fluctuations associated with that dissipation.

The point is, if there is a path by which energy can leave this system and go into the environment, then there is a path by which energy from the environment can come back into the system. The energy leaving the plate might be ordered energy as in the case of acoustic radiation. But energy can come back along the same path -- disordered energy -- and excite the plate to motion. That is the essence of the theorem: if you have dissipation, you have a connection to the outside world, and the outside world shakes your system, putting energy back in.

You reach an equilibrium state, where there is some level of energy going away from the plate structure and some level of energy coming back in. One might be ordered, one might be disordered, they both might be disordered. At any rate, you reach equilibrium and you can predict the level of vibrations in that structure.

(Transparency: Equilibrium-Thermal Fluctuations)

I am going to talk about statistical quantities. I am going to talk about quantities that are average quantities, and I will try to point those out, where they appear. I am trying to stay away from pathological cases so I do not get wrapped up in unimportant difficulties. We are talking

about systems with lots of molecules -- systems not supercooled or anything exotic. I am just talking about everyday sensors and about average quantities.

In that context, equipartition says that each degree of freedom of a system has a thermal energy of one-half Boltzmann's constant times the absolute temperature. A number of people have talked about this and you ought to pay attention to things that many people talk about, because they are either important or they are bandwagons. Since this is not a sponsored project review, we are probably talking about important physics and not bandwagons.

Picture a particle, a molecule in a liquid, for example: the molecule has several degrees of freedom. A number of people have talked about this. It can have kinetic energy in this direction, kinetic energy in this direction, and kinetic energy in this direction (three orthogonal directions).

Each of those energies, $\frac{1}{2} m v_x^2$, gets, on average, $\frac{1}{2} k_B T$. You can either take a snapshot of the system and average a bunch of molecules and get this number or you can follow one molecule through many, many, many collisions and get this number. So they are average numbers but, on average, the kinetic energy in any one direction is $\frac{1}{2} kT$.

The thing that a lot of people omit, once they get to this point, is how global this degree-of-freedom idea is. It is not only a molecule in a liquid that has $\frac{1}{2} m v_x^2$ equal to $\frac{1}{2} kT$, but if you suspend a ball bearing in the liquid it shows the same relationship. In the second case, m is the mass of the ball bearing; in the first case, m is the mass of the molecule.

You can see that the velocity of the ball bearing is going to be quite a bit smaller than the velocity of the molecule, but it has a velocity that you can calculate. Equipartition does not apply to just molecules. It applies to pink superballs, too. (Dr. Gabrielson is holding a pink superball suspended by a thin elastic band.)

This is a damped mass-spring system and it has several different modes of vibration. The most prominent is this up-and-down oscillation. It has kinetic energy and there is also spring-potential energy associated with the motion.

If you looked at just the up-and-down motion, you would have two degrees of freedom. One would give a kinetic-like term; one would give a potential-like term. You could either figure out the thermal-limiting value for the amplitude-squared (the AVERAGE amplitude squared) or you could figure out the thermal limiting value for the velocity-squared for this system.

STUDENT: Does this assumption prove the theorem, the equipartition theorem, for the ball bearing or this ball?

DR. GABRIELSON: I am not going to talk about proofs for this. There is some interesting literature that is quite old that gives plausibility arguments, and so on, for this kind of a system. Some of the arguments are pretty involved, but there is a lot of older literature on that stuff.

DR. MIGLIORI: If you put two springs in parallel, that is still only one degree of freedom. It is really the potential, not the number of springs.

DR. GABRIELSON: Yes, good point. Do not be misled by what the pictures are saying. It is the expression for this energy term that makes the difference. If I put two springs on the system, then this mode still has the same relationship. However, with two springs you can have a mode like this, where one spring is contracting and one spring is expanding. That gives you another degree of freedom and that gives you another contribution. Do not get too hung up on the simplicity of this picture.

Let us go back to the superball. Are those the only two modes? No. There is one like this (a pendulum motion). There is one where the ball twists back and forth; it has a very low resonance frequency. But beyond these modes associated the elastic band, the ball itself has modes. It has a breathing mode and it has a number of other modes.

The thing that saves you in doing a sensor calculation is that there is usually one mode that dominates the problem. If I made an accelerometer out of the superball-on-a-string, for example, by putting a pointer on the ball and a fixed pointer here on my body, I can measure acceleration with this system. The dominant mode in that case is this one: up and down. The other ones do not contribute very much. You can calculate them and make sure that they do not. But every system has a lot of modes and many are not important in a particular application.

You can demonstrate a lot of interesting physics with the superball-on-a-string. You can see the difference in phasing in the drive of an oscillator. If I drive it one way, I can get the amplitudes to grow. If I drive it another way, I can get the amplitudes to get pretty small. You can feel the drive phasing in your fingers. This is one of those things you really should do yourself.

You can also demonstrate that vibration isolation works. (Dr. Gabrielson is moving the end of the elastic up and down as fast as he can with as large an amplitude as he can while the ball on the other end of the elastic is hardly moving at all.)

(Laughter)

(Transparency: Energy Levels)

Some background about energy levels: there are lots of ways of describing energy. I am just going to write a couple of common energy levels in electron volts, which are 1.6×10^{-19} joules. It does not really matter what number I pick; they just come out to be nice numbers when I pick electron volts.

A number that you should all come to know and love is the thermal energy, Boltzmann's constant times absolute temperature, at room temperature. That is about 25 milli-eV. Compare that to the energy of a photon at optical frequencies, which is a couple of electron volts. A lot more energetic than a thermal agitation oscillation.

Next, consider an acoustic wave with a pressure of 100 μPa in air. That is a fairly low level. The energy density is 0.4 eV/cm^3 . That would be pretty small on the scale of a molecule. An acoustic wave with the same pressure level in water also has quite a small level -- $30 \mu\text{eV/cm}^3$. Very small when you get down to the size scale of a molecule.

But interestingly enough, compare this to steady state sonoluminescence and see the energy compression in that process. If you calculate the energy from the million photons coming out of each flash, each photon having a few eV and compare that to the energy in the acoustic field in the volume of the bubble, the photon energy is much higher than the acoustic energy. So in steady state sonoluminescence you can see the concentration of energy from a diffuse acoustic field to this very, very high energy (of very short duration) associated with the production of light.

Just for fun, take a look at chemical bonds. Typical chemical bonds require from 2 to 4 eV to break. Thermally you do not break too many chemical bonds but, optically, you do. That leads to photosynthesis and my personal favorite, sunburn.

Hydrogen bonds, which play a great role in the physical properties of water and also play a great role in biochemistry -- the way our enzymes and proteins function -- break at about 0.2 eV, much closer to the thermal level. At 0.2 eV, we are moving up the exponential tail of the distribution for thermal energy, so a whole lot more hydrogen bonds are going to be affected than other kinds of bonds. A lot of interesting chemistry here.

(Transparency: Sensing Gravity by Protozoa)

A lot of interesting biology also. I do some work in microfabricated sensors -- kind of a fascinating business where you try to make things smaller and smaller. There are lots of really interesting stories that come out of that and people continue to develop smaller and smaller devices.

This protozoan, which is not too big to start with, has this tiny little area up near -- I was going to say its head, but I suppose that would be presumptuous. Although maybe not, because it is mostly hollow.

(Laughter)

Inside this cavity, which is about $7 \mu\text{m}$ in diameter, there is a barium-sulfate granule attached to a hair-like structure. As near as anyone can tell, this thing acts as a pendulous accelerometer. It is really more of a trigger than an accelerometer in that when the organism is facing up, the trigger is pulled, and when the organism is facing down, it is not pulled. This generates a chemical response that is fed to the cilia along its body. It needs to know which way up is and which way down is in order to find the optimum oxygen concentrations in which it likes to live.

I think this is probably the smallest mass-produced accelerometer that I have ever seen. The micromachined-sensor industry does not like this, though, because it reproduces. You sell one and the business crashes.

(Laughter)

What can it sense? If you take a look at the sensor's potential energy, the ability to change potential energy inside the system with this proof mass and the allowable range of mass motion, and take the ratio of that to the thermal-agitation energy, you get about 300.

So we are starting to approach the level where you are going to have trouble sensing gravity. If you shrink the size of the organism, in linear dimension by four, this ratio goes to one and the organism is in big trouble. So the organism has some room for growth and some room for accreting the barium-sulfate granule in the first place.

DR. GARRETT: Bob Keolian says that that organism is the smartest organism on the planet per unit volume.

(Laughter)

DR. GABRIELSON: That is easy to believe. I would like to talk about the strategy that it uses to find the optimum oxygen concentration -- I do not have time for that here -- but it is really pretty fascinating.

(Transparency: Sensing of Magnetic Fields...)

Not to be outdone, there is a bacterium that also figures out which way up is. It uses a much simpler strategy: all it wants to do is get back down into the sediments at the bottoms of ponds. That is where its food is. But if the ponds get roiled up and the sediment is disturbed, it gets out into the water and it wants to get back down. There is nothing nastier than a perturbed bacterium, so this one has a mechanism for getting back down into the sediment.

The bacterium actually accumulates particles of magnetite. The particles are so small that they are single-domain ferromagnetic materials, which means you do not have to magnetize them, they already are: they line up very nicely by themselves. There is no intelligence associated with this mechanism; this is just a compass needle with goo stuck on. (Its mother probably would not like that description, but it is good enough.)

The dipole moment of these bacteria has been measured. The magnetic flux density for the Earth's magnetic field is roughly 50 microtesla. Figure out the potential energy in the magnetic field, compare it to the thermal energy and you get a ratio of 16. It is a smaller ratio than the protozoan, which may be indicative of the fact that its strategy for navigation is a lot simpler or it may be a consequence of the fact that the protozoan has an additional chemical transduction step to activate the cilia.

But the bacterium has to have some overhead, because it reproduces by splitting and giving four grains to each of the daughters and then they each accumulate a little bit more.

It is interesting to think about some of the popular science-fiction ideas these days, where they are talking tiny little micro- or nano-creatures that run around and assemble things.

If you figure out the world in which a creature on the scale of a nanometer lives, it is far different from what is described in these stories. The world of molecular collision dominates at the 1-nm scale. At that scale, anything is going to be batted around like crazy by molecular agitation.

If you had this picture in your mind of these creatures finding a place to go, moving in, grabbing on, performing some function, and then moving off to the right and down Broad Street and up First Avenue, that is not what is going to happen. These things are being bounced around all over the place because of molecular agitation.

This is one mechanism in the chemical world to move molecules around -- diffusion -- but it is certainly not a good mechanism for homing in and bringing one creature in contact with another one. This is not in your notes, but I did some quick calculations based on random walk, which is not really relevant to the kinds of sensors we are talking about.

If you took a body that is one meter in diameter and the same density as water -- that is the physicist's human, right, a one-meter sphere of water? The displacement in one second because of molecular agitation is about a nanometer. I cannot sense that, really. If somebody asks, "Why did you do that?" "Oh, molecular agitation." is not a good answer.

(Laughter)

I am sure that argument is going to be used in court some day.

You go to a millimeter with the same type of body and its displacement in one second is 20 nm -- still not that big a deal. At one micrometer, the body has almost a micrometer displacement in one second, so now it is going about a body length in a second. At that level, at the micrometer level, the body is going to be substantially affected by molecular agitation.

Let us take our nano-swimmer, a one-nanometer sphere neutrally buoyant in water. How far does it go in a second? It goes 20 μm , which 20,000 body lengths in a second. Picture that! Picture trying to do some intelligent navigation in that environment. So much for science fiction bashing.

(Transparency: Equilibrium-Thermal Fluctuations)

This idea of equipartition is a very powerful thing in thinking about systems and getting a feel for how systems work, but it is not all that useful when you are trying to describe the detailed behavior of a system from the point of view of building a sensor.

What is often more important is the distribution of that energy with frequency. Fortunately, a couple of people, Nyquist, notably, figured out in general what that distribution with frequency should be. I have written it for a mechanical case and for an electrical case.

F_n is force. This is the mean-square value of the fluctuating force. It is equal to four times Boltzmann's constant times absolute temperature times the mechanical resistance times whatever increment of bandwidth you are looking at. Mechanical resistance is force over velocity.

You can also do this with pressure and then the resistance would be pressure over volume velocity. You just have to pick the right units. You could make this voltage: the resistance is then voltage over current.

This is the one (voltage- or Johnson-noise) most people see with great regularity. This is a perfectly good analog of mechanical noise. In fact, I will make the case later that Johnson noise is fundamentally mechanical in nature.

This is not the whole story. Actually, the R is the real part of the IMPEDANCE. You can have a complex impedance: you take the real part and you have the right answer. The real part can be a function of frequency, in which case you get some shape to the distribution in frequency.

I would like to explain one thing about notation that confuses many people. If you have a noise quantity -- let us say it is a fluctuating noise force -- the noise is distributed over some range of frequencies: it does not have all its energy concentrated at a few discrete sinusoidal components.

To describe this, you describe it as a density, the energy or power in a certain band that you can measure. You can write this in a number of different ways. You can specify the bandwidth and write something squared. Let us do it for velocity noise. If we are talking about kinetic energy, we can talk about velocity squared.

So you could talk about the velocity squared in a 1 Hz bandwidth or in a 10 Hz bandwidth and so on. But you can normalize the bandwidth out and you can talk about velocity squared PER HERTZ. That is acceptable, too, because that is giving you a density.

A lot of people, because they like to speak in terms of measurable physical quantities, take the square root of the quantity and use terms like velocity per SQUARE ROOT of hertz. So do not be confused when you see an expression like this; do not be bothered by the fact that you are taking the square root of a frequency. You are actually talking about the distribution of power or energy in the frequency domain. It is an artifact of wanting to go back into a more natural unit.

DR. MIGLIORI: Are you going to talk about where these (expressions for thermal-noise density) roll off?

DR. GABRIELSON: Only in a crude sense: I have given the low-frequency approximation for Nyquist's theorem.

DR. GARRETT: But not very low.

DR. GABRIELSON: No, not very low. At EXTREMELY high frequency, the power density rolls off because of quantization. I am not going to talk about that, because not too many of my sensors get up that high, but there is plenty of literature about that.

DR. MIGLIORI: Yes, that is what prevents it from violating one of those thermodynamics laws.

DR. GABRIELSON: Yes, good point. If you integrate the density expressions I gave you over all frequency, you get an infinite amount of energy. So you know that is not right. Fortunately, not all of us like to integrate, so we do not get into that problem and we do not discover new physical laws.

You have to be very, very careful when you are talking about noise not to forget the fact that you really want to DO something with this sensor. I know I made light of that before, but you should not forget it. I made the joke the other night that I can reduce the noise of any sensor just by soldering a piece of wire from one terminal to the other. The noise goes way down, it is a tremendous improvement in noise performance, but the signal goes away too.

One way to avoid that problem is to talk in terms of noise-equivalent signal. In other words, do not give the noise in terms of the voltage output of the transducer; give the noise in terms of what the equivalent sensing signal would be to produce that noise level. That is one way of not getting into this problem (there are other ways of avoiding the problem, too).

DR. MIGLIORI: If you solder a wire across the sensor, you do not change the noise power. You change the noise voltage. The voltage goes down and the noise current goes up. Thermodynamics strikes again.

DR. GABRIELSON: Yes. Normally, though, you put this thing into a preamplifier and you measure the voltage.

DR. MIGLIORI: I am just picking here.

DR. GABRIELSON: It is a good point. Nobody who will pay me to solder a wire across the terminals understands the argument at that level.

(Laughter)

But yes, you do not get rid of the power. You change it into a nice current going through that wire.

(Transparency: Noise Equivalent Signal)

Take something like this (the accelerometer on the transparency). By now you know how to draw the proper circuit for this mechanical structure. This is an accelerometer. It is a mass-spring system in a case. You are measuring this relative motion (z). And, you have a noise fluctuating force (F_n).

Set the noise (F_n) to zero and solve for the signal response (z). How much z do you get for moving the case? Then set the signal (a) to zero and solve for the output (z) due to the noise. Then calculate the noise-equivalent signal.

It is a pretty simple process. In some cases it is even simpler than that.

(Transparency: Noise Equivalent Signal)

I will give you the answers for the two easy ones that I showed you pictures of on the previous transparency. I am not going to go through the whole process that I described. It is instructive to do what I said for these simple sensors, because it is relatively easy, but I do not even have to do that much work in this case.

M times the noise-equivalent acceleration is the noise-equivalent force on the proof mass. Set that equal to the fluctuation force -- the thermal agitation -- and you get the right answer. The noise-equivalent acceleration for this accelerometer is equal to $4 k_B T R / m^2$ which is also equal to $4 k_B T$ times the resonance frequency (2π times the resonance frequency, that is) divided by the mass, divided by the Q .

If you grind out the details of this problem, you will discover that the noise by itself shows the resonance peak and so does the signal by itself. The signal-to-noise RATIO, however, is flat (independent of frequency).

How do you make this sensor better if the thermal noise is what limits the sensor noise? Change Boltzmann's constant!

(Laughter)

Seriously, changing the temperature is a legitimate thing to do. You can do that. In my business, that is not a really smart thing to do -- well, in my previous business. We used to make sensors to throw out of airplanes into the water and we wanted to make the sensors small.

I can make a very, very small sensor perform well if I put it in a liquid-helium Dewar, but that does not make much sense for an inexpensive, off-the-shelf, expendable probe. But yes, you can lower the noise by lowering the temperature.

You can improve the signal-to-noise ratio by lowering the resonance frequency. You can increase the mass. Now, notice, if you are deciding between reducing the temperature and increasing the mass, in most cases I would prefer to increase the mass by a factor of four rather than to put it in liquid nitrogen (which would reduce the absolute temperature by a factor of four).

DR. MAYNARD: It (the noise-equivalent signal power) actually goes like $1/m^2$ (according to the $4k_B T R / m^2$ form) which gives an even greater advantage to increasing the mass.

DR. GABRIELSON: I am inclined to argue with that a little bit, depending on the particular damping mechanism. In some applications, you would be right while in others the R has an implicit dependence on the mass. For example, if the damping is squeeze-film damping

(viscous loss associated with squeezing a gas or liquid between two plates), changing the mass of the moving plate might be done by changing the area of the plate, which, in turn, changes the damping. I understand what you are saying, but I do not think it is always right.

It is a good point, though. If you have a system with FIXED damping, then you get a stronger effect by changing the mass than by changing the temperature.

You can also increase the Q of the sensor to improve the signal-to-noise ratio. The typical sensor philosophy is to make the sensor response flat over the band of interest and then roll the response off above that.

A typical sensor has its primary resonance at the upper end (sometimes at the lower end) of the band with critical damping so that the response is flat over the band. But nowadays, why sacrifice noise performance just to have flat response at the sensor? Why, with all the electronics we have? It is far easier to microfabricate electronic signal processing than it is to microfabricate a sensor.

What can you do? You can make the Q of the sensor resonance higher. B&K make a microphone that does that. Also, you can move the resonance into the band. If you put it in the band you may have to flatten the response electronically but it will work just fine, and give you better noise performance.

The same ideas apply to the pressure sensor. If you know the fundamentals, you can come up with sound designs.

(Transparency: Frequency Distribution of Noise)

You do not get rid of the noise power, though. In one sense, you can only squish it around. For any system, you ought to be able to integrate the noise output over all frequencies and you had better get the same answer as the equipartition prediction by doing that.

I have plotted some examples. These are two different sensors. They differ in their Q 's. If you integrate under the curves, you get the same total power for each sensor. You have to be a little bit careful to talk about appropriate quantities: the y-axis must be linear power and the x-axis must be linear frequency. Also, if you are calculating the noise-equivalent velocity, for example, the two sensors must have the same mass because equipartition only guarantees that the total kinetic energy be the same.

(Transparency: Noise Associated with Radiation)

I have a neat example that I would like to go through with you, because it shows that you can do the analysis two ways that seem to be very, very different and wind up with the same answer. This is what is enjoyable about working with equilibrium thermal noise. Some people have even demonstrated that you can get the physics wrong and still get the right answer. Not recommended though.

Let us talk about the fluctuation noise associated with acoustic radiation. It is not often a dominant problem at lower frequencies, but it can dominate at higher frequencies. Take a look at a spherical transducer, just a tiny little spherical transducer.

It has a spherical wave associated with it and you can write that like this -- you have seen expressions like this before. The pressure is going as $1/r$ and it has a complex-exponential dependence.

Use Newton's law to calculate the radial particle velocity. This is a straightforward differentiation of the first equation. Now you can write the mechanical radiation resistance, which is the ratio of force to velocity. The force is the pressure times the area of the sphere right at the sphere surface.

The impedance is a complex number. What part is going to tell you about the noise? The real part. So you take the real part of this impedance, and we are going to do it for a small source, so we will let kr go to zero. (This is not necessary: it just simplifies the discussion.) You get this expression: the density of the water times the sound speed in the water times the surface area of that source times the quantity, wave number -- ω/C -- times radius of the source, squared. You can simplify that a little to be π times the density times the frequency squared divided by the sound speed times the area squared.

You can then write the fluctuation pressure. It would have been a little bit simpler to use the pressure form of impedance: pressure over volume velocity. The final result is four times kT -- you had better know what that is by now -- times $\pi \rho f^2 c$ df.

(Transparency: Noise Associated with Radiation)

Now, there is another way you can solve this problem. There are two things going on here. Remember, I said that if there is a path away from the transducer into the environment, then there ought to be a path from the environment back into the transducer. Let us calculate the forcing of the environment on the transducer.

We will start out by considering a point receiver, now, centered in a cube, where the walls are a distance, L , away from the center. The acoustic modes for pressure -- let us make the walls rigid -- are going to look like sine waves with maxima at the walls. They are going to have forms like this: cosine of an integer, l , times $\pi x/L$; cosine of another integer, m , times $\pi y/L$; and cosine of another integer, n , times $\pi z/L$. This tells you what the modes are.

The wave numbers associated with those modes are: k_x equal to l times π divided by L , and the same for k_y and k_z , except using n and m .

We can plot the points of those wave numbers in wave-number space. Let us construct a figure. (This is a tool for solving this problem. There are other ways of doing this.) Plot all these points on a k diagram, where k_x goes this way, k_y goes this way, k_z goes this way. You get a cubic lattice of points if you plot all the points for all the integers, l , m , and n .

They have a definite spacing between them. The spacing between any two of them is π/L . If you make L very large -- the distance between the walls very, very large -- then the spacing (in k -space) between modes becomes very, very small and we can treat this as a continuum of points in k , which is what we are going to do. The actual wave number, of course, is the square root of the sum of the squares of k_x , k_y , and k_z .

Each cell in k -space has a volume of $(\pi/L)^3$ and there is one k -space point per cell. Each mode has a kinetic degree-of-freedom and a potential degree-of-freedom, so each mode gets a total thermal energy of kT . That means that in this k -space representation the density of thermal energy looks like this -- the energy per mode times the modal density in k -space. Isn't this so much easier than the last method?

(Transparency: Noise Associated with Radiation)

I really want the distribution in frequency, though. I do not want the distribution in k . The first thing I will do is figure out how many modes are between k and k plus δk . That is just a geometry problem: find the volume in a spherical shell.

You get energy density in terms of dk and now I can translate it into frequency because $k = 2\pi f/c$ and $dk = 2\pi df/c$. I now have an expression for the energy density in frequency. Divide that by the spatial volume and that gives the energy density in terms of space and the distribution with frequency.

You know from some of the other talks that another way to write the energy density is p -squared-RMS over ρc -squared. If you compare these two expressions, you can solve for p -squared-RMS and, wonder of wonders, you get the same answer that you got from the radiation resistance.

In one sense this amounts to no more than another derivation of Nyquist's result, but it is kind of neat, because the approaches are so different.

(Transparency: The Simple Accelerometer)

I am going to go quickly over the next few transparencies, because it is really a sample problem for you to look at later to see if you can wrestle with it on your own. I am looking at the simple accelerometer and actually treating it as a velocity sensor, because I have a measurement of a system like this. Write an equivalent circuit for it, solve for the signal response of the device, and solve for the noise response.

(Transparency: Equilibrium Noise in a Geophone)

The important thing to notice when you are working your way through this is that the equivalent voltage output of this transducer for the mechanically associated noise has a shape something like this (curve 1). It peaks at the resonance frequency.

There is another component that is very important in this case and that is the Johnson noise associated with the resistance of the coil winding. (The measurement example I am going to use

is for a geophone.) It gives another thermal-noise expression and, over this frequency range, the noise is white (curve 2). Again, I will let you read through that, but I wanted to show you a measurement of geophone noise.

(Transparency: Noise-Floor Measurement)

This is measured noise for a geophone (it is actually a geophone array, but that does not make any difference). The predicted Johnson noise just from the coil winding resistance is at this level (where the PREDICTED curve intersects the left scale), so there is a white noise component and that is the Johnson noise associated with the electrical resistance of the winding.

Added to that is the mechanical-agitation noise -- the thermal equilibrium noise associated with the moving mass. That is pretty fair agreement between theory and measurement. This is not an easy measurement to make, because these acceleration levels are about 10 BILLIONTHS of a g per root hertz: a lot of isolation is required to make this kind of measurement. But once you remove all the extraneous pickup, you get a pretty nice comparison with the prediction.

DR. GARRETT: What are those spikes?

DR. GABRIELSON: The two tallest peaks are resonances in the second stage of the vibration-isolation system. The inner chamber is suspended from a tripod and the tripod is a little too flimsy -- I want to rebuild it. The tripod has a twisting resonance and a bowing resonance in the legs. The peak at 5 Hz is a resonance in the third stage of the vibration isolation. That is a very crude isolation, just a lead block on a couple of rubber bands inside the inner chamber.

I do not know where the other smaller peaks are, but they seem to appear in many of my lab experiments. Maybe they will shift when I get up to Penn State.

(Laughter)

(Transparency: Equivalent Noise Generators)

If you like to draw equivalent circuits of mechanical systems, I will point out that it is very easy to account for thermal-equilibrium noise. The analogy between mechanical circuits -- force and velocity -- and electrical circuits -- voltage and current -- not only holds for the differential equations governing the dynamics, it also holds for the thermodynamics. If you put these electrical generators in, you get the right answer for the equilibrium-thermal noise.

(Transparency: Signal-to-Noise Ratio: A Useful...)

There is a slick little theorem that I found in a book. I have not had the guts to go through the proof myself -- well, actually, I do not care. It seems to work pretty well. When you are analyzing relatively complicated circuits, this can save you a lot of time.

The theorem says that the signal-to-noise ratio at the output of a linear circuit does not depend on the value of the output load. So what? Let us say you have this circuit, which is moderately complex. There is a signal on the input, noise on the input, a little noise added here,

some kind of shaping gain (A) here to the load, feedback (B), and additional noise in the feedback path.

Calculation of signal-to-noise ratio for even a simple circuit like this is awkward because of the feedback loop. But if the signal-to-noise ratio is independent of the output load, make the output load zero -- actually, you can do it two ways, make the output load very, very small and then, when you are done, make it go to zero, or you can work with output currents.

If you calculate the output-current signal-to-noise ratio, you get the right answer. But by shorting out the output, you open up the feedback loop. This makes the calculation a lot easier. You do not have this feedback path to worry about. I have tried it on a number of circuits and it works just fine.

(Transparency: Signal-to-Noise Ratio and Feedback)

With this technique you can address another interesting issue. Can you improve the noise of a system by adding feedback to it? You can add feedback to a mechanical system and, by doing so, you can effectively create negative resistance in the system.

If we add negative resistance to the mechanical system it should drop the noise, right? Let us just take it from the point of view of this circuit. Here is the system we want to change (upper left). There is some output signal-to-noise ratio from this system. Let us put feedback on it and we get some other output signal-to-noise ratio (upper right).

We can use the theorem, short out the output, and calculate the output-current signal-to-noise ratio. If you take a look at this circuit (bottom left), the feedback path (B) can only ADD noise to this circuit. It cannot reduce it. You always get poorer performance in a feedback circuit than you do without feedback, although sometimes the increment is so small or the benefits of feedback are so great that you should use the feedback anyway.

The point is that the signal-to-noise ratio is not improved this way. Why doesn't it work? Because the negative resistance produced by the feedback is not a damping mechanism that is in equilibrium with the environment. Sure, I can get negative resistance in the circuit, but it is not an EQUILIBRIUM resistance any more.

Any fluctuating voltage is being fed back with complete coherence so it is not in equilibrium with the external environment and you cannot draw conclusions as if it were. A useful thing to do, though, with feedback is to design a system with very, very low damping -- a large Q -- and then put feedback on it to make the response more manageable.

The equilibrium thermal noise is set by the EQUILIBRIUM damping, which, in that case, is small. You wind up with some signal processing problems by doing that, so you cannot carry the process too far. But it will work, because the true resistance from which the equilibrium noise is being calculated is not changed.

DR. MIGLIORI: If the feedback generator is the same as the amplifier, so that you have coherent noise, and the feedback impedance is reactive, do you get any noise cancellation at that point? You do not increase the noise if you are feeding back, but do you get a square root of 2 or something like that come in?

DR. GABRIELSON: Yes, you might.

DR. MIGLIORI: I do not really understand that problem very well.

DR. GABRIELSON: I do not, either. If you do the analysis with this particular theorem, though, you might take a little bit closer look at it.

DR. MIGLIORI: The thing you would worry about is, in proving this theorem, are the noise sources correlated or uncorrelated in a feedback supply. That is the only way you could possibly win.

DR. GABRIELSON: Yes. In everything that I have talked about so far I have assumed that all of these noise sources are uncorrelated, and that may not be true. They may be correlated through some process, so you do have to be a little bit careful with that.

The analysis is easy enough to modify to handle coherence between noise sources. Just be careful: if you grind through these calculations, you get results that are very comparable to measurements, if you make the measurements carefully enough.

DR. MAYNARD: There is a real example of this. Pat Arnott (Desert Research Institute) was doing flow acoustics where a light signal (a modulated laser) was exciting the resonator. He was using a thermoacoustic prime mover to enhance the Q of the resonator to improve the signal-to-noise.

DR. GABRIELSON: I have used his system to illustrate the signal-to-noise improvement problem in some other talks. He called me when he was working on that problem and I worked out the theory for the noise and the modes in the resonator. You get only the improvement you hope for up to the point where you reach the thermal-noise limit and then it stops. You cannot go beyond the thermal-noise limit with that setup. You get rid of OTHER noise problems by increasing the signal output, but only until you reach the thermal-noise limit. It is a nice problem.

(Transparency: Quiz Question #2)

I put this in as a quiz question and Professor Marston demonstrated to me that he is conscientious enough to have read the notes ahead of time. Several days ago he told me that this question might not be probing exactly what I wanted to probe. He was right.

You should actually replace the voltage source with a current source. The effect that I wanted to show is present in either circuit but, unfortunately, there is a complicating effect in the one I drew.

The complicating effect is that the internal impedance of an ideal voltage source is zero. So when you close the switch, you are actually shorting this resistor and you are taking that thermal noise contribution out of the problem, whereas with the current source you are not: the internal impedance of an ideal current source is infinite.

What I wanted to get at, though, is are you adding noise because you are putting a current of random carriers through the resistor. Random carriers arrive randomly and so produce shot noise, and I wanted you to think about how much the shot noise increases the total noise in the system.

(Transparency: Shot Noise)

First, let us take a look at what shot noise is. Sometimes it is considered to be a source of noise separate and distinct from thermal-equilibrium noise. Sometimes it is, sometimes it is not. If you take a process that produces impulses and the impulses are randomly spaced, you have a picture like this. I drew this as an electrical current, but the events could be molecular collisions, too.

The current is a sum of delta functions and with amplitude, Q , which is the electrical charge, in this case. Expand the current as a Fourier series. These are the kinds of integrals that you love to get. If there is a delta function inside the integral, great. Expand it as a Fourier series, get the coefficients in the usual way, and because the delta function is here, the integral is simple: it turns into a sum.

The period over which you are integrating really does not matter; just make it long enough so that there are lots of events in there. The period does determine the effective analysis band.

Take one component of that series and compute the mean-square value. The cross term -- cosine time sine -- averages to zero. You are left with these two terms. Each one time averages to one-half so you get half of the two squares of those coefficients.

(Transparency: Shot Noise)

If you take the expression for the a_k coefficient, square it, and time average it, you again get some terms that are products of equal cosines -- the argument is the same -- and then you get all the cross products.

How do the factors in the cross products differ? By the time indices. In the first term you are multiplying two factors with identical time indices together. In the second term, you are multiplying factors with two different time indices together. In SOME circumstances, the second term (after time averaging) goes to zero. For this to happen these time intervals must be truly independent of each other and random. (In other words, the events themselves must be independent.) In that case, the time average of the cross-product goes to zero.

If you work out the power in the k th component, then the mean-square fluctuation current is $2q$ times the mean (DC) current times Δf . This is the classic shot-noise expression. A lot

of people go through this in terms of Poisson's statistics and that is a nice way to do it, too, but this shows you that you can do it another way. It actually gives you some advantages when the events are not delta functions; it is a little easier to see this way.

But remember, for this expression to be true, you must have impulse-like disturbances and they must be independent.

(Transparency: Circuit for Basic Noise Experiment)

So let us do the experiment. This is a fun little experiment to do -- and easy. This is a low-noise amplifier. I would not do it this way, again (there are better circuits to use). But it is a good low-noise amplifier. Normally there would be one resistor from the gate of the FET to ground and the input would be applied to the gate.

If one gate resistor is good, four gate resistors must be great. You replace the normal bias resistor by this configuration of resistors. If the battery is not here at all, what is the equivalent resistance hanging on the gate? One megohm.

If you short the battery and the switch, what is the equivalent resistance on the gate? One megohm.

So the resistance the gate sees does not depend on the characteristics of the battery, as long as the bridge is balanced properly. You can get away with putting almost anything there where the battery is. This is actually why I did not stumble onto the problem in my quiz question: because I did the actual experiment with the bridge.

Now, I did not put the bridge in for that reason -- I am not that smart. I did it because I wanted to cancel any noise introduced by the battery. If the battery introduces noise, that noise cancels out in the legs of the bridge.

I also wanted to do it this way because I did not want to change the dc level at the gate (which would change the bias point and the circuit gain). At any rate, it worked out just fine.

You make the measurement with this circuit and what do you get?

(Transparency: Resistor Noise)

Here is the output of a spectrum analyzer: the flat wiggly line is the voltage noise when the switch is open -- just the noise of a one-megohm resistor. Calculate the Johnson noise for a one-megohm resistor and it is right here. This is the sort of agreement you ought to get when making noise measurements.

I put a variety of batteries in to try this circuit out. I went all the way up to 90 V. Our lab stockroom still gives out 45-V batteries for the B-plus supply for vacuum-tube circuits.

(Laughter)

The wiggly curves that rise at the low-frequency end are for just a couple of volts. If you close the switch, you put a current in the bridge. If you calculate the shot noise of the current

through the resistance and you get an equivalent voltage noise that it should be up here ("Shot Noise Predicted").

You get no measurable difference at the higher frequencies whether the switch is open or closed, in spite of the fact that the shot noise ought to be substantially higher. The rise in noise at the low-frequency end is, by the way, $1/f$ noise (I am not going to talk about that much but it is very important in practical systems).

At any rate, you do this experiment and you get no significant difference with and without current.

(Transparency: Noise from Molecular Collisions)

Let us do a problem in molecular collisions to see if the situation is any better. Maybe there is something strange about electrons. Let us do a problem in molecular collisions. Look at this little piston. The force exerted on this piston -- I am going to look at collisions on just one side of the piston -- you can write as the rate of change in momentum of a molecule initially traveling to the right, hitting the disk, and bouncing back. Sum over all molecules and you get the total force.

Another way of writing that is the flux of molecules -- the time rate through a unit area -- times the momentum change per collision. N is the number density of the molecules (molecules per cubic meter), A is the area of the face of the piston. The average of the x -component of the velocity with which the molecules are arriving at the piston is v_x . Notice that this quantity in parentheses is the average x -component of the molecular velocity minus the velocity with which the disk is trying to get out of the way. The momentum change is $2mv$ (it has mv coming in and it has mv going out, so it must have transferred $2mv$).

Work this product out and you get three interesting terms. The third term is proportional to the speed at which the disk is moving, squared. I want to do the problem where the disk speed is small compared to the molecular speed.

The first term, if you look in your kinetic theory books, looks like a pressure on the diaphragm -- pressure times area. That makes sense. The middle term is a factor times the disk velocity, x -dot: a factor multiplying the disk velocity is a damping factor. So we have a mechanical resistance equal to two times the number density of molecules times the mass of the molecules times the average velocity in the x direction times the area of the disk.

Now that you know what the mechanical damping is, you should be able to write an expression for the noise, right? Everybody say "right", even if you do not believe it.

("Right")

Thank you.

(Transparency: Noise from Molecular Collisions)

The fluctuation force squared is $4 kT R_{\text{mech}} df$. I want to work in terms of pressure so I include the area and get this expression. Again, kinetics tells me that the mean pressure, P_0 , is nkT .

Also, the mean velocity is twice the mean of any of the components. Look that up in any kinetic-theory book.

So the fluctuation pressure equals 2 times the quantity carried by the carrier -- this is the amount of something that can be transferred in a collision -- times pressure times df divided by area. It looks like a shot-noise expression.

Actually, it looks even more like a shot-noise expression than it appears at first glance.

(Transparency: Generalized Forms for Shot Noise)

The general form for a shot-noise expression looks like this: the mean-square fluctuation in some flux density (of anything) is equal to two times the amount that is carried by the carrier (you have to have a carrier if you have discrete events; if you do not have discrete events, we cannot talk about shot noise) times the average flux density times the bandwidth divided by the area.

For electric charge -- I do not know if you ever see it written this way -- the fluctuation-current density (squared) is equal to two times Q (Q for each carrier) times the average current density times Δf over A . Multiply this expression by A squared and you get the standard shot-noise form: i -squared is equal to $2qI_0 df$.

But it is nice to write it as shown on the transparency because it allows you to write fluctuation expressions for other processes a bit easier. For example, for photons you can get the right expression immediately. The mean-square fluctuation in optical intensity is 2 times hf times the mean intensity times Δf over A . How about momentum flux density: what is momentum flux density? It is pressure. The mean-square fluctuation pressure is equal to 2 times the momentum transferred per carrier times the mean pressure times Δf over A . All of the expressions look alike.

If you have a process where the events are discrete and independent, you can write the mean-square fluctuation this way. Notice that this fluctuation pressure is proportional to the mean pressure. It is actually a higher power than linear in pressure, because the momentum is also a function of pressure. So this expression is a function of pressure, a strong function of pressure.

(Transparency: Molecular-Impact Noise)

We can handle molecular-impact noise in another way. We can do it from a hydrodynamic point of view; we do not have to do this molecular nonsense. Hydrodynamics gives you the right macroscopic physics and we are dealing with equilibrium thermodynamics so we still ought to get the right answer.

We need the mechanical resistance. For a disk moving broadside, Stokes gives a nice expression for the mechanical resistance for small-velocity motions: 16 times the viscosity times the radius of the disk. The fluctuation-pressure squared is then $4kT$ times 16 times viscosity times the radius of the disk divided by the disk area.

This expression is INDEPENDENT of the mean pressure, roughly speaking. (Viscosity is a very weak function of pressure.) Uh oh, there is a problem!

Well, I tricked you. I said that shot noise requires that the collisions, or events, be not only discrete but also be independent. However, the hydrodynamic solution applies to the case where the mean-free-path between collisions is very, very small compared to the size of the disk.

The collisions are anything but independent in that case. A molecule out here has no chance at all of hitting this disk undisturbed. It is going to hit many other molecules on the way in. You know it gets there eventually, but it has no chance of getting in undisturbed from here.

A molecule that has just bounced off the disk has no chance at all of getting out here undisturbed. The molecules do not act independently. In fact, the act of moving this disk disturbs the local velocity distribution in the molecules, and that is a very important effect if you are trying to do this calculation microscopically.

Now, you would be crazy to do the calculation microscopically for the case where the mean-free-path is small compared to the disk. I am crazy, but I am not sufficiently talented to express my craziness in that particular way.

(Laughter)

We have actually analyzed two different problems. If you noticed, at the top of the first slide of the molecular-collision discussion were the words "free-molecular flow." The analysis that I did for you with the molecular collisions is perfectly valid if the mean-free-path is LARGE compared to the disk, when the collisions are independent events.

(Transparency: Molecular-Impact Noise)

If you take the ratio of the two expressions, you get $3\pi/8$ -- which is one, for small values of π -- times the radius of the disk over the mean-free-path. This is the controlling nondimensional parameter for determining which expression is correct.

When the mean-free-path is large compared to the radius, you are down here in this region (left of the dashed line on the graph) where the noise fluctuations depend strongly on pressure. If the mean-free-path is small with respect to the radius, you are up here (right of the dashed line) where there is not much dependence on mean pressure.

(Transparency: Noise in Metallic Conductors)

What about electron flow? We ought to be able to make some sense out of that. I am giving you quite a bit of background in these next few transparencies that I am not going to talk about, and it is really for your information, to be able to look at this problem in different ways.

(Transparency: Noise in Semiconductors)

I will just point out that I am as sloppy as the next guy -- each of the expressions for mean-square fluctuation should be followed by a Δf for bandwidth. But let us skip over these details, and try to understand the difference between a metal and a semiconductor.

(Transparency: $k_x - k_y$ Grid)

Do you remember when I was talking about the radiation problem for sound, I said that you could plot the modes on a k -space diagram? You can do the same thing with the electromagnetics problem in a cube of metal. You get wavefunctions that fit in the cube -- lots and lots of them -- and you can fill up k -space with those points.

In a metal there are plenty of electrons in the conduction band, lots and lots of electrons. No two electrons can occupy the same state -- the same spin-state, that is. They fill up different states, and as you have more and more electrons in the conduction band, k -space fills up from the origin outwards farther and farther (from the lowest energy state upwards).

If you plot the modes in k -space for a metal, k -space looks like this. (Dr. Gabrielson puts a disk on the transparency centered on the origin.) There are lots and lots of points and they fill up a circular region (a sphere in 3-D). Most materials do not produce a perfectly spherical region, but for the purposes of this argument we will look at it this way.

This picture is the picture of a metal at absolute zero -- or at least very low temperature. There is a very sharp dividing line between the filled states and the unfilled states. The surface is very well-defined. And the radius of the disk is pretty large. In copper, for example -- the surface is 7 eV. Compare 7 eV to the 25 meV that you get from thermal agitation and you will see what happens when you warm the metal up.

When you warm the metal to room temperature, it looks like this. (Dr. Gabrielson raises the disk up above the transparency a few inches on a pedestal so the projected image is that of a disk with a fuzzy edge.) The space is still full on the inside, but now it looks fuzzy on the surface. The last time I did this, someone came up to me afterward and said, "I didn't get it. I saw a fuzzy circle the first time and a fuzzy circle the second time."

(Laughter)

At any rate, this is what it looks like at room temperature. The thermal agitations have affected this edge and made this edge fuzzy, but if there is an electron in a state well inside the disk, it would have to go several electron volts to find another open state. There are no open states nearby, so those "deep" electrons cannot fluctuate. Only the electrons in this little fuzzy layer at the edge of the sphere have vacant states near enough to allow them to fluctuate. That is fundamental to understanding the noise production.

If you put a voltage on this metal, what happens? The whole sphere in k -space shifts a little bit (disk moved a small fraction of a diameter to the right). The shift that I showed is a substantial voltage, but it is just a very small shift away from the center.

Now there is a net imbalance in the momentum vectors which results in the measurable current. But notice that there is very little difference in the pattern of fluctuation -- the pattern of the fuzzy surface -- with respect to the origin of the diagram. Fluctuations do NOT add coherently so there is no measurable difference between the observed fluctuations with and without applied voltage. Consequently, you cannot use shot theory based simply on the mean current for current fluctuations in a metal. In a metal you measure Johnson noise from zero DC-current to very high currents. To get any increase over Johnson noise, you would have to shift the fuzzy surface a significant fraction of a radius away from the origin, but this requires so much current that the metal may melt before a substantial deviation from Johnson noise would be noticed.

How about a semiconductor? This is a picture of a semiconductor at absolute zero (the k_x - k_y grid with the disk removed). There are no conduction electrons. (I am talking about a normal intrinsic semiconductor -- nothing pathological.) In a semiconductor EVERY conduction electron is in the conduction band because it was thermally excited into that band.

If you raise the temperature, this is what it looks like (Dr. Gabrielson holds a small ball a few inches above the center of the grid to project a small, fuzzy dot). It looks like a little fuzzy dot at the origin: a quantum dust bunny. There is no completely filled region. If you put a voltage on the semiconductor, the dot moves away from the origin. All of the electrons that participate in the conduction also participate in the fluctuations. So in a semiconductor you do get true shot noise -- and, perhaps, noise from other mechanisms.

If you understand the k -space diagram, you understand most of what is happening in conduction. I am going to put these next figures up just to make absolutely sure that you are confused, but I want a chance to show you what the equations say about the conduction problem.

(Transparency: Occupation (f) and Fluctuation (df))

This is the well-known occupancy function for electrons (equation in lower left). This (f) is the occupancy. The occupancy is one if the state is occupied at all times. In a normal metal, this is what the occupancy function looks like as a function of energy. The states are fully occupied up to pretty high energies and then there is this little band, df , where thermal agitation has its effect. This small band is responsible for the agitation.

(Transparency: Occupation (f) and Fluctuation (df))

In a semiconductor you can use the same occupancy function; however, E_f , which is the Fermi level, is way down here in the semiconductor (well below the zero reference on the eV

scale). When you work out that same equation -- it goes to a limit as an exponential -- you get a very small occupancy: states in the conduction band are only infrequently occupied.

Every electron has neighboring states into which it can fluctuate. In fact, df and f are identical in this case, in stark contrast to the metallic case. And, there are not many carriers in the conduction band.

(Transparency: Noise as a Function of Applied...)

I included the next two slides just to give you a problem to work through. It is an interesting little puzzle. This came up when I was working in electron-tunneling sensors. In electron tunneling there is a current that can flow between two conductors that are spaced very close together (but not touching). The current flow is extremely sensitive to the spacing so a number of interesting measurement applications use electron tunneling.

Tunneling is quantum mechanical, so you would expect that the events are independent and you would expect that there is a shot noise associated with it. But you can also measure the resistance of this junction. You put a voltage on, measure the current, and calculate a resistance. Some people have argued that there is also a thermal noise associated with the resistance and that the thermal noise should be added to the shot noise to get the total noise.

But that is not what happens. In the case of a tunnel junction, for example, there is a forward current and a reverse current. The reverse current goes uphill, the forward current goes downhill through this potential barrier. They are related by a Boltzmann distribution with this exponential form (upper right). The measured current is the difference between these currents while the measured current fluctuation is the square root of the sum of the squares of the individual current-fluctuation components.

(Transparency: Noise as a Function of Applied...)

If you do the calculation, you get a very interesting and satisfying result. If the voltage you put across this system -- ΔV -- is such that $q \Delta V$ is much, much smaller than the thermal energy, kT , the expression degenerates into Johnson noise (expressed in terms of current noise).

If the voltage across the system is much, much greater than the thermal energy, the same equation degenerates into shot noise, so it is very nice. The real noise is a function of applied voltage and goes like this (the solid line in the graph). It is not the sum of a separate shot noise component added to a separate Johnson noise component. Everything is right with the world. It all works out very nicely.

I am going to skip the avalanche photodetector (transparency: Photodetector Shot Noise) and talk about shot noise if the generating events are not impulses.

(Transparency: Shot Noise)

Real, physical events are often quite different than ideal impulses. If the events have some shape, then you can model them as a delta-function input to some system with the appropriate transfer function.

All I need to do to modify the theory, then, is to modify the result by the Fourier transform of the impulse response of the generating system.

You take this string of impulses and feed it through the system with some characteristic response and out comes the desired waveform.

(Transparency: Shot Noise (Single...))

What if you pick something like an exponentially decaying response? This is a simple thing to do because the exponential is easy to integrate. The impulse response is an exponential decay with some time constant.

Now, instead of the noise being white out to the quantum limit, it is shaped like this, with an f^2 dependence (in power) beyond this characteristic frequency.

(Transparency: Shot Noise (Multiple...))

If you have lots of processes, though, and those processes all have different characteristic times, you can get a complicated frequency dependence. Some people have shown that for many, many distribution functions having different characteristic times, you get something that is pretty close to a $1/f$ power law in the decay region.

This may be the way $1/f$ distributions are produced, but it may not be. It is plausible but not necessarily correct.

(Transparency: $1/f$ Noise)

And $1/f$ noise is important, because it is all over the place: in physical systems, in electrical systems, in chemical systems. There are many very interesting papers about $1/f$ distributions. I am not going to talk about it too much, except to make some key observations to keep you from going too far astray in the literature.

If you take a system that is in equilibrium, you can calculate the thermal-equilibrium noise for that system. The noise is not going to be any different than the equilibrium calculation. If you pump energy into the system, though, it may depart from that equilibrium calculation. If it is a very nice metal and you pump current through it, it is not going to depart measurably, but a lot of other systems will: systems with poor junctions or surfaces, semiconductors, many other things.

You put power into them and their noise grows above the equilibrium. Often those noise components have distributions like $1/f$ instead of being white or $1/f^2$.

This $1/f$ behavior can extend over many decades in frequency. Now, integrating a pure $1/f$ spectrum down to zero frequency gives you infinite power, so that cannot be right. There has to be some point at low frequency where the spectrum stops rising. Measurements have been made

and they just run out of time. If you are making a measurement at a microhertz, it takes a megasecond (about 12 days) to do the measurement. It is hard to go decades below that.

So there must actually be a lower limit to the $1/f$ behavior. If this function integrated to a finite number, it would probably be easier to figure out what the underlying noise mechanism is, but it does not. So, to characterize the noise, you need not only the total noise power, you also need the low-frequency roll-off point.

I would like to close the discussion of noise with a short anecdote. Because of my involvement with miniaturization, I have seen a lot of proposals for very small sensors. A number of years ago, it struck me how few people calculated the basic equilibrium thermal limit for their proposed sensors. The sensors are small and so are affected much more than larger sensors. One proposer constructed an elaborate electrical equivalent circuit of a sensor and he used SPICE, which is a circuit simulator program, to calculate the response. Now, SPICE has a noise option in it and all he would have had to do was turn on the noise option and he would have had the thermal noise output. (Remember that. It is a useful technique.)

So I thought that I would write a short tutorial paper for one of the journals related to microsensors. I wrote a little paper explaining how to do these kinds of calculations for thermal equilibrium noise in very small sensors. It was trivial, almost embarrassing to write something like that.

It came back from review a while later with the comment, "If you're going to introduce new physics, you ought to have experiments to back it up."

(Laughter)

Needless to say, it was simple to find experimental results; I used some measurements from the early 1900's.

(Laughter)

But noise issues really are important and there are some excellent books around. Professor Maynard asked me to mention Kittel's book on statistical physics, a very nice treatment of many of these noise issues.

(Transparency: Reciprocity)

I am going to shift gears now and talk about some calibration techniques for sensors. You can learn a lot about a sensor by studying its self-noise. In fact, if you study it well enough, you can learn some things about the sensor that perhaps you would miss by studying only the signal response.

On the other hand, by trying to measure the response of the sensor as well as you possibly can, you may also uncover details of the sensor operation that you would not find otherwise. Precision calibration techniques are good for that. They will drive you crazy, but if you keep

tweaking and tweaking and getting better and better at your measurement of the sensor response, you will uncover things that you previously did not appreciate about the sensor operation.

I am going to talk briefly about two different techniques for calibrating sensors. The first is reciprocity. Here we are going to treat the sensor or device as a black box with two ports. I have drawn this one with electrical terminals on one side and mechanical terminals on the other side -- but it could be electrical on both sides, it could be mechanical on both sides, and so on. There is a potential quantity, e , and a flow quantity, i , on one side, and a potential quantity, f , and a flow quantity, v , on the other side.

Do not make a big deal about the physical significance of potential and flow. All I am doing is giving you notation. There is no problem switching these quantities around. None of the electrical engineers I know like this, but you can interchange the roles of e and i . You can make voltage the flow quantity and current the potential quantity, and, as long as you are consistent, you get the right answer. What we are doing is constructing a framework that obeys the proper differential equations.

Now, e can be a function of the input current and it can be a function of the input velocity. The quantity, f , can be a function of the input current and it can be a function of the input velocity. If the system is linear and you can write a set of equations like this for this two-port network, frequently, but not always, $B = C$.

If that is true, then the device is reciprocal and you can play some wonderful games. If $B = C$, then you can write this kind of relationship: the voltage for a given velocity input with the current set to zero -- which is B -- equals the force for a given current input, with the velocity equal to zero -- which is C .

Notice that you never have to say the word "reciprocity." The third equation follows from the first two. Reciprocity is just a label for this kind of system. It is not something magic.

Many kinds of transducers are reciprocal, many are not. The ones listed on the left often are reciprocity to relatively high degrees, at least for small signals. There are some transducers that are decidedly nonreciprocal: piezo-resistive sensors, for example. Put some resistive material on a bar and bend the bar. The resistive material is strained and the resistance changes.

Apply a voltage to it, measure the change in resistance, and relate that change to the change in stress. You have made a transducer. However, if you put a current through that resistor, you might get some bending of the bar from thermal expansion, but you do not get reciprocal equations for the system. Electron-tunneling transduction is also not reciprocal.

DR. RASPET: B and C , are those just numbers, or are those functions, because it seems that if you have linear equations and those are constants, you can always make $B = C$.

DR. GABRIELSON: If B and C are simple, non-zero constants, then yes, the equations can be scaled to make B and C equal. If B or C is zero, then no; if B and C are functions of

frequency, temperature, humidity, whatever, and the functional dependencies are different, then, again, you cannot scale to construct a reciprocal system.

(Transparency: Reciprocity)

In some systems $B = -C$. There is a whole body of literature addressing this, but it is unnecessary. If you have a system with $B = -C$, you can always, in this kind of equation set, rearrange variables so that some new B equals some new C .

In the previous example, I had e and f on the left sides of the equations. Now I have rearranged the equations to have e and v on the left. If I do the transformation, the new B equals the new C . You never have to talk about anti-reciprocal networks. You just rearrange things so that they are reciprocal.

I will say that $B = C$ does not always guarantee the system to be reciprocal: if $B = C = 0$, then the two ports are not coupled.

(Transparency: Reciprocity)

Let us take a look at a simple electrical circuit. I am trying to explain reciprocity from the same perspective in which I learned it, because I had some serious areas of misunderstanding about reciprocity. It is the 90's so I will blame it on my parents.

(Laughter)

I am going to try to show you my blunders to prevent you from falling in the same holes.

Here is a network. It is electrical on both sides. It has two resistors in it -- easy to analyze. You see that e_1 is a function of i_1 and i_2 , and e_2 is a function of i_1 and i_2 . Lo and behold, $B = C$. It looks like it should be reciprocal. Notice that these coefficients look like transfer impedances: a voltage on one terminal divided by a current on the other terminal.

First major point: reciprocity does not depend on the device being lossless. This is not a lossless device. There is plenty of dissipation in this resistive circuit.

I was visiting another lab and I was talking to the people in the calibration group. They do not use reciprocity, because they believe that a device must be lossless for reciprocity to be true. Wrong.

Second major point: reciprocity does not mean that VOLTAGE ratios are identical. You have to take the right quantities to get equal ratios. If you put 10 V into this circuit and get 1 V out and then you turn it around and put 1 V in, you do not get either 10 V or 0.1 V here. Reciprocity does not say anything about the voltage ratios. It says something about the transfer impedances and nothing else. You have to play the game by the rules.

(Transparency: Reciprocity Calibration)

Consider some transducer. I am going to start just with definitions. Alpha is the receiving response. These are arbitrary definitions, but I am picking them to rig the problem, as any good instructor would. Alpha equals the voltage output of the sensor for a given velocity input. I am

putting some velocity into the sensor and measuring the open-circuit voltage out. That is my definition of receiving response (not appropriate to an accelerometer but appropriate to a geophone).

Transmitting response I call beta. If I put a current into this device, it produces a force. The force output per current input is beta. You can pick all kinds of definitions here and you can still do reciprocity just fine; these are just convenient.

Why calibrate by reciprocity? Because it is often difficult to make mechanical measurements accurately. It is often difficult to measure a velocity or a force directly, so you would like some way around that. What you can do is rearrange the problem so that the difficult measurements are electrical and the easy measurements are mechanical.

DR. ATCHLEY: In your diagram, does f become input (flow) and v become potential?

DR. GABRIELSON: Yes. I did that just to keep your minds open to the possibilities. This is what you would get with the classical "anti-reciprocal transducer." You would flip the function of the quantities on one terminal thereby making it reciprocal.

With a geophone, for example, you have to do this to get a truly reciprocal transducer. It is important to realize that this is a mathematical construction to solve the problem.

I could have swapped quantities at the other terminals; I could have left f and v alone and switched i and e . Then I would have found out who the electrical engineers were, because they would have run, screaming, for the door.

(Transparency: Reciprocity Calibration)

The classic way to do reciprocity is with three transducers and three measurements. You need a source transducer (which does not have to be reciprocal); a receiving transducer (which also does not have to be reciprocal); and at least one reciprocal transducer.

You connect pairs of transducers through some transfer impedance. This is always true, although the transfer impedance is not always obvious. For example, for two hydrophones in water the transfer impedance is produced by the water medium. In a standing wave tube, the transfer impedance is produced by the medium confined in the wave tube, and so on. You can always identify a transfer impedance, though, and you can always cast the classical results this way.

Drive this system, SZR, and measure the output. Drive this system, SZT, measure the output, and drive this system, TZR, and measure the output. Form these ratios. The capital letters in square brackets represent these ratio measurements.

There are three equations. How many unknowns are there? One here, the receiving response of R. One here, the transmitting response of S. And two here, the receiving response and the transmitting response of T. So I have three equations and four unknowns. I guess I have to pay a consultant to finish the problem.

(Transparency: Reciprocity Calibration)

This is where reciprocity comes in. You get the fourth equation from the reciprocal relationship for that T transducer. Here are the measurements cast in terms of the responses and the transfer impedance.

If you multiply SR by TR and divide by ST, the result is $1/Z$ times $\alpha-R$ squared times $\beta-T$ over $\alpha-T$. If I can figure out the relationship between $\beta-T$ and $\alpha-T$, I am done, because then I would have the receiving response of R. For different combinations you can get different responses.

The transfer impedance does not have to be the same for each measurement. You would just get a Z_3 over $Z_1 Z_2$ factor instead of the $1/Z$. But they do not have to be the same. In a lot of the classical derivations they are the same and some times it is easy to make them the same in an experiment.

How do you get the relationship that you want between β and α ? In what is probably the most beautiful paper on reciprocity, Professor Rudnick gives a procedure for doing this by using a hypothetical transducer on one end and the reciprocal transducer on the other end.

He has some very nice ideas in that paper. I distributed it. If you work through that paper from the front to the back, you will know more about reciprocity than most people who use reciprocity on a daily basis.

You can do a very similar sort of thing here. You connect this transducer, in your mind, to some hypothetical reciprocal transfer impedance. Since both T and Z_0 are reciprocal, the combination is reciprocal and you can write this relationship to get the result that $\alpha-T$ equals $\beta-T$. Now, that result depends on the choice for α and β . You will always get a relationship between these two, but there might be other factors in there if you chose another form for α and β .

The problem is solved. In practical terms, what you need to do is figure out some clever way of putting in a transfer impedance that is easy to measure, because you need the properties of the transfer impedance. You need this Z. That is part of the measurement.

There is a paper on reciprocity with an absolutely ugly title: something like "Calibration of Transducers Without an Absolute Standard." I know what the author means, and I do not mean to quibble, but that is baloney. You always have an absolute measurement associated with a calibration. What he means is you are not comparing the test transducer against an absolute standard TRANSDUCER. You ARE making absolute measurements, though, and I will show you some examples later.

(Transparency: Reciprocity Calibration (variations))

There are lots of variations possible in reciprocity calibration. If you measure the currents electrically by measuring the voltage across some resistor through which the current is flowing,

then you can actually use voltage ratios rather than absolute voltages. Voltage ratios and resistances are easier to measure in my lab than absolute voltages.

Another interesting thing to do is to set up the drive values so that you reproduce the same field at the output transducer in pairs of the measurements. That is really nice, because if you have nonlinear S or R, you can still do the calibration: you run the nonlinear transducer at the same point each time. That is covered in Professor Rudnick's paper.

If you have two reciprocal transducers, then you can explicitly check to see whether reciprocity works in your apparatus. You can run a current into the first, and measure the voltage out of the second; then run a current into the second and measure the voltage out of the first to see if the ratio is the same over the whole frequency range that you want to run. It is a nice test.

If you have two identical reciprocal transducers, you only have to make two measurements. Of course, you will never get the uncertainty between the two this way. Sometimes you can use one transducer; you can excite a lightly damped system and watch the ringing, or you can get reflection off of, say, an air-water interface.

(Transparency: Reciprocity Calibration (examples))

Here are some examples of situations in which you can do reciprocity calibration. The free-field calibration is probably the most frequently used example to illustrate reciprocity. One transducer is emitting sound and the other transducer is receiving it. There are no nearby boundaries to produce reflections or standing waves.

The transducers are at a certain separation. You use S and R, then S and T, then T and R in this same environment, and you get all the measurements you need.

Do you need to make an absolute measurement of anything? Of course. The transfer impedance is ρ times the frequency over $2\pi r$. You need to measure r absolutely and you need to measure the density of the water (fortunately, that is not too difficult) and, of course, frequency.

You can do more, though. Whenever range is a parameter in a free-field measurement, one thing to do is to make several sets of measurements at different ranges, then you can least-squares fit the results over those ranges and you can reduce some errors.

For example, measuring the absolute distance between the transducers is not straightforward: there are problems associated with location of the effective acoustic centers. But if you put the transducers on a fixture and change the relative spacing, the separation CHANGE can be measured more accurately than the absolute separation. So moving things makes a lot of sense in a case like this.

(Transparency: Reciprocity Calibration (examples))

There are two situations with very simple transfer impedances: This (upper diagram) is the way I calibrate geophones and accelerometers in the lab. Hang them on a mass suspended as a pendulum and make the measurements. The transfer impedance is the mechanical impedance of the mass -- $j\omega m$, or minus $i\omega m$, depending on the convention you select for the assumed complex-exponential time dependence.

There are many corrections you can make to reduce the error. You account for the dynamics of the geophones and the dynamics of the pendulum to make the results quite accurate.

But, again, you are making an absolute measurement: you are making an absolute measurement of mass. Fortunately, that is pretty easy. I would also point out that you must make an absolute measurement of the resistor that you are using to measure the current (if you are measuring voltage across a resistor to get current).

In a traveling wave tube the transfer impedance is simple: ρc over A (cross-sectional area of the tube). All you need to do is use the reciprocity relationships and put in the proper transfer impedance.

(Transparency: Reciprocity Calibration)

A really nice problem is the "rigid-walled" resonator. It is not a resonator with infinite Q however; it is a real resonator that might have a high but finite Q .

This takes a little bit more work to figure out. One way of doing it is to look at the stored potential and kinetic energies to get the total energy. Then calculate the energy lost per cycle. Then you put those quantities into the Q expression: 2π times the energy stored over the energy lost per cycle. Solve for the transfer impedance, which is now in terms of relatively easy-to-measure quantities. This problem is presented in Prof. Rudnick's paper.

Now the pressure is a maximum at both ends of the resonator to a very good approximation. But zero is NOT a good approximation for the velocity at the driving end. That is where the energy to sustain the oscillation is being supplied. The velocity cannot be zero or there would be no power input.

(Transparency: Reciprocity Calibration (Equip))

The pendulum setup is nice because it is easy to build and you can get good results with a minimum of effort. You can put all transducers on one fixture like this; you just change electrical connections to select transducer pairs.

Another fringe benefit of this is that you can put a mirror on the end of the pendulum and set up an optical interferometer to check the results. It is a very good idea to make measurements in at least two different ways.

DR. GARRETT: I would also like to point out one of the nice things about your pendulum calibration is that all three (geophones) really are reciprocal and, therefore, you are triply redundant in your determination. Even though you labeled them S, T, and R, it is T, T, and T.

You just treat one as S, one as R, and the other as T, and then you can treat the two that are together as S and R and the right-hand one as T. There is no particular reason, if they are all reciprocal, not to calibrate each one three different ways.

DR. GABRIELSON: Absolutely, do as much as you can with them. You can put a non-reciprocal transducer in this fixture and calibrate it, but I have been doing it a lot with geophones and piezo-ceramic accelerometers. There is a lot of flexibility there.

(Transparency: Reciprocity Calibration (Equip))

Here is a picture of one version of the interconnections just to show you how you might plug something like this together. I have taken the drive (sinusoidal or wideband noise) from the analyzer. You can also use a lock-in amplifier to do this; I just drew it for a spectrum analyzer.

I measure the voltage across this resistor (R_0) to get the current. I put R_0 in the ground lead, so that the source ground and the channel-one ground can be common. Measure the output of the receiving transducer in channel two.

Then you switch transducer connections as need be. It is a very simple setup.

(Transparency: Free-Mass Reciprocity Cal)

This is an example of such a calibration for a commercial piezoceramic accelerometer. The manufacturer provides a calibration for this accelerometer and the response values are given to five significant figures; their data is about 2% high compared to these measurements.

(Laughter)

I called the manufacturer and asked what the uncertainty was in the quoted calibration values. They said, "Well, if you want a value between two of the listed frequencies, you just interpolate between the two." I said, "No, no, that's not what I mean. Let's say you're right on one of the values that you gave me. What's the uncertainty, what's the plus or minus?"

He said, "Well, that's the number we measure. The number is on the sheet." I gave up. You realize right away that this concept of uncertainty is foreign, especially to someone who quotes a five-significant-figure number in a situation like this.

(Laughter)

They also publish a noise floor for this transducer but they cannot measure it. When I measured it, the noise floor was quite a bit higher than their specs.

DR. KEOLIAN: We found that some companies do not know what noise floor means. It is easier to be smarter than the professionals on this.

DR. GABRIELSON: Yes, but you also need to keep in mind that they are not getting paid to sit around and scratch their heads about these things, so they do not often take the incentive to do it on their own.

DR. KEOLIAN: They hired somebody to figure all that stuff out and then they fired him.

(Laughter)

DR. GABRIELSON: This graph is a combination of several results. I sent the unit to NIST to have them calibrate it. I did not tell them to do anything special, just an ordinary 1% calibration.

I did a reciprocity calibration for both a vertical pendulum and a horizontal pendulum. (The vertical "pendulum" was a mass suspended by surgical tubing.) I also did a Bessel-null calibration, which I am going to talk about next. There is reasonable agreement here. Plus or minus 1 dB is roughly from top to bottom of the graph.

I do not put dB scales on the sheet, by the way, because if I put this on the sheet, that the entire displayed range was plus or minus 1 dB, my boss would have told me to stop a long time ago. He would have said, "2 dB is good enough for us." I often wanted to ask him if it would be okay if I changed his salary by a dB.

(Laughter)

Then somebody objected and said, "You can't do that: for dB you need either an energy quantity or a power quantity." I said, "I don't see the problem. Money is a power."

(Laughter)

(Transparency: Bessel-Null Calibration)

There is another really valuable measurement technique that I cannot resist telling you about. Let us say we want to measure the motion of some driver or some vibrating surface. Let us take a laser -- now there are lots of ways of setting up an interferometer, but this is a quick and easy way -- and a fiber-coupler (a 3 dB coupler). One end of the coupler goes to the laser, one end goes to the photodetector, one end is placed near the moving surface, and one end is not connected.

Let us look at a blown-up view of the fiber end near the moving surface. The light is coming down the fiber, out of the fiber, hitting the surface and going back in: d is the displacement between the surface and the end of the fiber. Another component hits the end of the fiber and reflects back along the fiber from there. These two paths interfere and you measure the interference pattern in the photodetector. Straightforward.

DR. MIGLIORI: A tiny comment. If you are using a visible laser diode and you do not shatter the non-connected end of the fiber, the reflection kills the gain and shuts the laser off, if it is a really good grade of laser.

DR. GABRIELSON: Yes. I use a helium-neon laser that is more resistant to such things but I also put a dab of cement (like Duco) on the not-connected end to reduce the reflections.

(Transparency: Bessel-Null Calibration)

The component that reflects from the moving mirror varies as the cosine of the quantity $\omega_0 t + 2k_0 d$. ω_0 is the laser frequency (times 2π) and k_0 is ω_0 over the speed of light.

The component that reflects from the cleaved end varies as the cosine of $\omega_0 t$. The photodetector output gives the intensity of that light -- it responds to photon flux. You get signal [1] plus signal [2] quantity squared (and time averaged). The output of the photodetector is very simple: $\cos(2k_0 d)$.

Let us take a look at sinusoidal motion of the mirror: there is some static displacement and some sinusoidal displacement about that equilibrium position.

(Transparency: Interferometer Transfer Function)

The most common interferometer strategy is the small-signal arrangement. You adjust d_0 so that $\cos(2k_0 d)$ is at a zero crossing. You can use a little piezo-stretcher in the fiber or you can even physically move the transducer if you have sufficiently fine control.

This (lower left) is the input displacement -- the AC displacement -- and you get out something that looks like this (upper right), a nice sine wave. This is the small-signal operation. It is nearly linear if you keep d_1 small enough.

The spectral output is a single line (lower right). To calibrate this arrangement, you can increase and decrease the displacement (statically) to locate the modulation peaks and derive the zero-crossing slope from that.

But if the system drifts, then this zero point drifts off to the side and the response changes. It is really a pain in the neck.

(Transparency: Interferometer Transfer Function)

However, if you drive the system harder (which seems to be a common theme in some of the talks here -- driving things harder to see what happens) the output goes strongly nonlinear. You are driving it to the modulation peak and further into subsequent valleys and peaks: you get this relatively complicated waveshape (upper right).

If you look at the spectrum (lower right), you see harmonics. I bet some of you can guess how the amplitudes of these harmonics are related by now, because at least two other people have talked about similar kinds of systems with similar harmonic structure. They are related by Bessel functions to the displacement amplitude.

(Transparency: Large-Signal Photodetector...)

A neat little measurement that Professor Garrett suggested to me is to see if you can generate a Bessel function through this measurement. It is fairly easy to obtain a nice Bessel function but I chose to do things the hard way. I was looking for an interferometer technique that would work on my reciprocity pendulum. (The mass is not clamped.) This pendulum has a Q of about 600 at a frequency of 0.7 Hz. It moves a lot more than an optical wavelength even without drive. So I wanted a technique that would work for that arrangement.

These data are taken from the pendulum, so there is a lot of noise in the curve. If you do it on a solidly fixed transducer and you are careful, you can "draw" a really nice Bessel function. The solid line, by the way, is a Bessel function that was just an eyeball fit to the data.

To generate this curve, I varied the drive amplitude at the frequency, f_1 . Here, I drove the pendulum at 100 Hz. I slowly increased the drive amplitude and I measured the output of a narrow-band detector at f_1 .

Increase the drive amplitude and the output goes up initially, levels off, comes back down, you get a nice null, further down, up, another nice null, and so on. The fact that you get those nulls is very nice.

(Transparency: Bessel-Null Calibration)

The photodetector output at the drive frequency is this $\sin(2k_0d_0)$ function times the Bessel function, $J_1(2k_0d_1)$. You can look up the values for the arguments that produce zeros in J_1 in many standard references.

So you know that $2k_0d_1$ at the first null is going to be 3.83171. If you are using a helium-neon laser, in particular, you know k_0 very well. That is certainly not limiting the uncertainty in my measurement. So you can calculate d_1 directly and that is what you wanted in the first place.

Not only is it simple to get accurate determinations of displacement, the process is not sensitive to drift away from some specific equilibrium point. The only way you get messed up is if you happen to wind up on a null of the sine of $2k_0d_0$. Then you get poor signal-to-noise ratio or no signal at all.

In my apparatus, the pendulum is moving very slowly over many optical wavelengths, so I never have to worry about it sitting at a null of the sine function. I just have to integrate long enough -- at least a second or so.

This is a neat technique to use on a moving pendulum. It would be extremely hard to make a small-signal measurement on the pendulum because of the oscillatory drift, but the large-signal method works very nicely. Also, you can get multiple nulls by increasing the drive further. Or, you can look at the second harmonic of the drive frequency and find the nulls of J_2 . You can go to the third harmonic and look at the nulls of J_3 . You can get a lot of points at one drive frequency.

Since each null is produced for a different drive level, you can tell something about the linearity of the system. There are a lot of clever things you can do with this technique: I wanted to mention it, because many people are not aware of how powerful it is as a technique for obtaining absolute displacement measurements. The technique provided a nice check on the reciprocity results.

(Transparency: Components of a Sensor System)

I would like to close by putting things in some context. Normally a transducer is part of a sensing system: knowing about the transducer itself is only part of the problem. You really need to look at the transducer in the system in which it is embedded. It does not make much sense to do a detailed calculation of thermal agitation noise for a sensor whose electronic noise floor is 10 times greater.

You need to know what the controlling values for noise are at each stage. There is noise in the transducer itself and that was most of my focus: the intrinsic noise in the transducer. But there is not much point in lowering the noise floor of the transducer many times lower than the ambient external noise floor in which the transducer will be operating. You ought to know what the expected ambient is in order to make a smart design.

Conventional noise is not the only issue. If the sensor is supposed to respond to pressure but it also responds strongly to temperature, that is something that you do not want. It is a legitimate transduction phenomenon, but it can degrade the performance of your system. This is "cross-selectivity noise" or "spurious response."

The preamplifier may limit the overall performance. Perhaps 90 percent of the literature on sensor noise says that the important thing is the preamplifier; perhaps 90 percent of the time that is right. But the other 10 percent of the time it is not right and absolutely ridiculous things are said about what you can do with a transducer.

So you ought to know the fundamentals. If you calculate the fundamental effects, then that will save you from making at least some stupid remarks.

Many people assume that the sensor system is only a transducer and a preamp, but there are many systems that are more complex. More and more frequently, systems are being designed with some kind of feedback. It may be an electromechanical balancing sort of feedback -- the electron tunneling transducer is a good example of that -- and the feedback is going to add noise as well.

You might do direct digitization in the sensing loop itself. That is a popular technique for surface-microfabricated sensors that have very, very small sense capacitances so the electronics must be integrated on the same chip. The easiest process to integrate is CMOS (complementary metal-oxide semiconductor). However, CMOS has a lot of inherent noise, particularly at low frequencies. What is normally done then is to design the sensor to be self-modulated so that the sensed signal is translated up out of the low-frequency noise from the CMOS. The modulation scheme can either produce an analog output or a digital output.

Transducers may require a carrier or a bias. Capacitive microphones usually have a DC bias voltage that can add noise. There are not many microphones that use AC bias, but you can get remarkable performance that way, especially at low frequencies. For one thing, AC bias gets you away from $1/f$ noise.

You have to look at the whole system and see what kinds of problems you are going to have. Even if you are doing a calibration and you get things you do not understand, they may not be in the transducer -- they may be in the bias or in the feedback or somewhere else.

From a practical standpoint, the issues of noise are very important, particularly if you are pushing technology. I think a lot of you will get into positions where you are going to try to push technology. You had better know about these things.

You had better know how to measure (and predict!) the response of your system, so it is good to look at a couple of different ways of doing that. Reciprocity is a lot of fun, because you can think of a new apparatus with some transfer impedance and make a different kind of a measurement.

When I was first fooling with reciprocity, I was doing reciprocity calculations of little electrical networks with electrical transfer impedances. It is a really silly thing to do, because you can measure the input of an electrical circuit and the output of an electrical circuit and get everything directly and easily.

But you can also do reciprocity measurements to convince yourself that you understand reciprocity. If you have the time, I am an advocate of doing silly things, because you learn that way. Unfortunately, I do silly things naturally. Thank you very much.

DR. MARSTON: Could you put up the transparency of the fiber interferometer, again?

(Transparency: Bessel-Null Cal -- 1st slide)

There is a measurement strategy that is being advocated by some people and that is best explained in this picture by removing the 3 dB coupler and replacing that by your laser, so that, in effect, the light comes back through the fiber, back in through the output coupler on the laser, and then you are measuring the light on the back end of the laser in the photodetector.

I recently went to a talk where somebody said this was a wonderful way to do things, and I asked him afterward if he could tell me in a simple way what this does to your noise and what it buys you. I did not get an answer.

I was wondering if anybody in this room could tell me what that buys me? You can answer the question, if you know.

DR. GABRIELSON: I do not know the answer to that.

DR. GARRETT: Ed Carome may be a better guy to answer this. But I will take it, if Dr. Carome does not want to take it. What happens when you have a laser is that the chip is a resonator, it has these two phases. When you let the light that is reflecting back from your mirror get back into the laser, what you do is you superimpose a second laser cavity that is extraordinarily long, which means it has very small spacing between modes.

When you superimpose that on your laser to some unknown degree, you do not know how much the coupling is to the external cavity as opposed to the internal cavity, and it just drives the

noise bonkers. It really makes your diodes noisy. In fiberoptic sensors we go to great lengths to try to make sure that the back reflection into the laser is small because of the noise.

DR. MARSTON: I still want to understand what is the beneficial effect? I can see lots of complications coming in due to the problem of saturation and so on, but what is the beneficial effect?

DR. GARRETT: I do not see any.

DR. MIGLIORI: What Professor Garrett says is correct. With a diode laser it is a disaster. First of all, the coherence length is shorter than the cavity, so all it does is introduce noise, which actually will shut the laser off at some point; that is where the big noise is coming from.

But the 3 dB coupler is also a noise source. If you remove the 3 dB coupler, you pick up a square root of 2 somewhere. You are just getting more of the light back into the system where you are measuring it.

Every time it goes through that 3 dB coupler, it splits in half. So that way you lose half, that way you lose half, and there is a whole bunch of 2's in there. That is the only reason for getting rid of it.

All of the other reasons totally overwhelm the reason for getting rid of the coupler, as far as I am concerned.

DR. CAROME: You are saying you used a helium-neon laser. That is important, because of the gap between the mirror and the fiber. You do not want to do this, I do not think, with a laser diode: the gap between the end of the fiber and the mirror is greater than the coherence length of the laser.

DR. GABRIELSON: You actually can get results with a gap spacing in excess of the coherence length of the diode laser by some small factor, but the helium-neon is easier to set up. Even I can get results from it.

DR. GARRETT: Plus if you put the tip of your fiber very close to the swinging pendulum, it is only a matter of time before the fiber and the pendulum become one.

DR. MIGLIORI: And then two.

DR. GARRETT: And then six.

(Laughter)

DR. HARGROVE: Let me remind you, however, that if you look at this hookup in a somewhat different way, the end of the fiber and the moving mirror looks like a scanning interferometer on the laser to measure the frequency of the modes, and so forth.

When I worked with helium-neon lasers in their infancy and you wanted to know what the mode structure was, you were always very careful to put an optical isolator in there so that this was not reflecting back into the cavity and making it go bonkers.

STUDENT: What size fiber do you use when you do this as far as the core? Do you have to use really small single-mode or can you use the larger?

DR. GABRIELSON: This is a single-mode fiber.

DR. GARRETT: There is signal to burn.

DR. GABRIELSON: Yes, this is not a minimum deductibility system by any means; I am nowhere near the noise floor.

DR. GARRETT: The nice thing about Bessel-null is that the result is signal-independent, because you are looking for zero crossing. In fact, your answer is zero signal.

DR. MIGLIORI: From looking at those dots (on the "measured" Bessel function), if you did not fit those dots to find the zero crossing, you are throwing away signal to noise.

DR. GABRIELSON: Yes, let me explain the technique as Professor Garrett explained it to me. This procedure, by the way, does not work well with the pendulum, but it will work with a fixed transducer.

Instead of trying to get a null, you set the drive at some amplitude away from the null a little bit and then average. If you are using a spectrum analyzer, you use peak hold, because you want to drift through the maximum response. You repeat the process for a number of drive levels bracketing the null. You fit those points with a line and get a nice zero crossing.

This is extremely difficult with the pendulum, because there is far too much sideband modulation from the Q-enhanced motion introduced at the 0.7 Hz pendulum resonance. (Since the modulation is nonlinear, the effects spill over into much higher frequencies.) So you do actually dial in the null in that system. For the pendulum, it is better to dial in each null many times and average the results. But yes, use all the information you have.

THERMOACOUSTIC ENGINES AND REFRIGERATORS

GREG SWIFT

LOS ALAMOS NATIONAL LABORATORY

DR. SWIFT: I want to start with a couple of viewgraphs that should be easy review for everybody, just to get us thinking about heat engines and refrigerators.

(Transparency)

You learned in your thermodynamics classes that there are two classes of heat engines: the heat engine, or prime mover, in which heat flow from a high-temperature source to a low-temperature sink produces work; and the refrigerator, or heat pump, in which the directions of all the energy flows are reversed, so that you put work into the device and lift heat from a low-temperature reservoir, depositing waste heat to a high-temperature reservoir.

(Transparency)

The first and second laws of thermodynamics put some bounds on the efficiencies of such devices. For the engine, the first law can be expressed like this (this is just energy conservation). The device itself is in steady state, so the energy going in must equal that going out.

The second law of thermodynamics is a little bit harder to explain in words. The second law says that the entropy of the universe always increases (or it can stay the same). The universe here is made of three pieces: the hot reservoir, the device, and the cold reservoir. The device is in long-term time-averaged steady state, so the only pieces of the universe we have to think about are the two reservoirs. The entropy decrease of this reservoir is Q_H divided by T_H , and the entropy increase of this reservoir is Q_C divided by T_C . So the second law says that the entropy increase of the cold reservoir has to be greater than or equal to the entropy lost from the hot reservoir.

(Transparency)

A little bit of algebra to combine these, eliminates Q_C to get an expression for the efficiency. What you want in an engine -- work -- divided by what you have to spend to get it -- Q_H -- is bounded by the Carnot efficiency, this ratio of temperatures.

For the refrigerator, you play the same games and do a little bit of algebra, eliminating the variable you are least interested in, in this case the waste heat rejection. You are interested in the so-called coefficient of performance, what you want -- cooling power -- divided by what you have to spend to get it -- work -- and it is bounded by this temperature ratio.

That should all be review.

(Transparency)

This should also be review for people in this audience -- a simple sound-wave derivation, a derivation of a Helmholtz equation. The notation I will be using throughout this talk is not much

different from what people previously have used. All the variables, such as pressure, will be written as a mean value plus the real part of a complex amplitude times e to the $i\omega t$.

In that approximation, the momentum equation, or Newton's law, ends up looking like this: force equals mass times acceleration. The continuity equation, which just says that if the density at a point is increasing, it must be because there is more stuff coming toward that point than going away from it, looks like this in the acoustic approximation.

The other piece of physics that is always needed for a wave-equation derivation is some kind of equation of state of the material, which we will express here as tying the oscillations in pressure to those of density by the adiabatic derivative of p with respect to ρ , which is the square of the sound speed.

I am sure you people have done this algebra many times. You substitute this into that and eliminate u and get a Helmholtz equation, like that, describing how p varies with x .

(Transparency)

It turns out that it was not until 1969 that anyone was able to correctly write down a similar equation for the following circumstance: parallel-sided channels -- ducts, if you will -- where lateral thermal contact between the sound wave and the duct walls is important, and with a temperature gradient in the duct.

Nicholas Rott published this in 1969. This is what we call his wave equation. You can see the pieces of the old familiar wave equation: 1 times p_1 is that, and here are the two derivatives of p_1 with respect to x . In the course of the next hour or so, we will look in much more detail at these new pieces: f_{nu} , the viscous correction terms that appear in a couple of places, and f_k , thermal correction terms.

Here we have a term that has just one derivative of p_1 with respect to x . It looks like it might be a damping term. It turns out it can be a gain term, as well as a damping term. That is one of the most important aspects of thermoacoustics.

(Transparency)

During the next few hours, we will be looking at a combined picture of heat engines and Rott's analysis of sound waves, forming the subject of thermoacoustic engines and refrigerators. Thermoacoustic engines and refrigerators generally have some kind of resonator. Imagine here a quarter-wavelength resonator, open at one end where sound is being radiated out into the surroundings. This is a heat engine: heat flows from a high-temperature source and heat flows into a low-temperature sink, producing work that is radiated into the surroundings.

Here is a schematic of the corresponding refrigerator, in which acoustic power is put in, in order to lift heat from a low-temperature reservoir and dump the waste heat to a high-temperature reservoir.

The next viewgraph shows some of the relevant length scales, just to get us oriented. Wavelength, of course, will be related to the length of the resonator, and the thermal penetration depth will be related to the gaps between things in the heat-exchange elements.

(Transparency)

Length scales: Along the propagation direction, an important length scale is the wavelength of the sound. Typically, we have half-wavelength resonators in thermoacoustics. The overall length of the apparatus will be some fraction of a wavelength.

Another important length scale in that same direction is the gas-displacement amplitude. If you mark one little parcel of gas and watch it move back and forth, how far does it move? That is often very large in thermoacoustic engines. This is drawn more or less to scale.

Perpendicular to the propagation direction, the relevant lengths are the diffusion lengths for heat and momentum. Heat diffuses. The question is, in an oscillating problem at frequency ω , in a time comparable to $1/\omega$, how far can heat diffuse? The answer is the thermal penetration depth, which turns out to be related to the thermal conductivity of the gas and its density and specific heat and the frequency. Similarly, how far can momentum diffuse laterally? The viscous penetration depth, related to the viscosity, density, and frequency.

Relative sizes of these penetration depths: If you do the algebra dividing the viscous penetration depth by the thermal penetration depth, you see that the ω s cancel out, so you get something that is frequency-independent. The quantity under the square root here, with viscosity on top and thermal conductivity below, is called the Prandtl number of a fluid.

For ideal gases, that is typically just a little bit less than 1. For monatomic gases, it is about 0.7, so the square root is about 0.85. If you make a monatomic gas mixture, such as helium with argon or helium with xenon, you can get this down to about a quarter.

In ordinary acoustics, such as that with which we are communicating, the gas-displacement amplitudes are quite small compared to the penetration depths, which in turn are quite small compared to the wavelengths, of course.

In thermoacoustics we turn this around a little bit. The gas-displacement amplitudes are usually quite a bit larger than the penetration depths and quite a bit smaller than the wavelength. Typical numbers might be a quarter of a millimeter for penetration depths, so the spacing here is likely to be a millimeter or thereabouts. Gas-displacement amplitudes can easily be centimeters, and the wavelength might be a meter or something like that.

(Transparency)

I have a couple of demonstrations that I will show you. The first is of a thermoacoustic engine. It is just like this picture. This is the little thermoacoustic engine. Between here and here is a quarter-wavelength resonator. It is closed on this end and open on this end. We will supply a heat source and a heat sink, and in a minute or so it will start to produce sound.

You can change the frequency quite easily by changing the geometry of the resonator, and it is perfectly happy with that. The stack length is about a centimeter. In this region there is a little stack of parallel plates, stainless steel. They are maybe 2 or 3 thousandths of an inch thick and separated by 10 thou or 15 thou, something like that. The heat exchangers are little copper strips with comparable spacing.

For a refrigerator demonstration, I have this: a half-wavelength resonator from here to here, filled with air. Down on this end is a little glass stack, with copper strips on the two ends to represent the heat exchangers. You can sort of see what that looks like. It is hard to see the thermocouple wires that sense the temperature difference between the two ends of the stack.

Parallel plates in the stack -- these are microscope cover slips. The electronics box is a piece of junk from Radio Shack. It used to be a public address system amplifier until we put in an oscillator and a bunch of other stuff. There is a differential thermocouple across the stack to measure the temperature difference. When I turn on the sound, this end of the stack will get hotter than ambient, this end will get colder, and the middle stays somewhere around ambient.

The trouble with this demonstration today is the following: Notice how raspy that sounds. It does not sound like a pure tone at all. That is because yesterday, when I was trying this out, I blew the loudspeaker out. Yesterday it got a little bit over a 20° C temperature difference. When I left Los Alamos, it would routinely go over a 60° C temperature difference. Mike Hayden, the postdoc who built this demonstration, warned me about driving it too hard!

(Transparency)

Some examples of engines and refrigerators, from the more or less real world, at least the world of applied research and development. This is a fun picture, because it shows three people wearing six hats.

(Laughter)

One of those people is David Gardner. This was a project to begin exploring thermoacoustics as a sonar projector. These unfortunate guys are at a lake in upstate New York in January a few years ago, and they are learning how hard it is to solder wires in -50° wind-chill weather. The thermoacoustic resonator was about 20 feet long, with stacks and heat exchangers and things underneath the electronics in the foreground. That is an example of a thermoacoustic engine.

(Transparency)

Here is another example. These are some people at Tektronix Corporation in Beaverton, Oregon. We teamed up with them to build a heat-driven cryogenic refrigerator. The thermoacoustic engine is everything from here to here, basically a half-wavelength resonator with two stacks, two hot heat exchangers, two room-temperature heat exchangers. We built that at Los

Alamos. That produces acoustic power from heat to drive this thing called an orifice pulse-tube refrigerator.

Here is an aside for the people in the back row so that they do not go to sleep. Most of this talk is about standing-wave thermoacoustic engines and refrigerators, but there are also traveling-wave thermoacoustic devices. The pulse-tube refrigerator is a sort of traveling-wave refrigerator, in that the essence of it is that a traveling wave goes upwards through the regenerator and heat-exchanger things, producing cooling at this end. This is also related to the work that Peter Ceperley was doing quite a few years ago. There just is not enough time in a three-hour talk to develop intuition about both the standing-wave engines and the traveling-wave engines, so we are going to pretty much ignore the traveling-wave stuff.

DR. LEVY: What temperatures could you achieve?

DR. SWIFT: This one got 150° K. We have done 90° K. Pulse-tube refrigerators that have been piston-driven rather than thermoacoustically driven have gotten to 4° K now. Four or five stages are required to do that. Fifty Kelvin is very routine in a single stage.

(Transparency)

On the refrigerator side of things, here is a thermoacoustic refrigerator built at Ford Motor Company. The meter stick shows the scale. It is a quarter-wave device from the top to here, with a large belly down here and a big woofer as a driver. The stack sits up in here, with hot heat exchanger here, cold heat exchanger here. You can see the pipes for cooling water. This was 10-bar helium. It ran about 400 cycles a second, maybe 50 watts of cooling power, something like that.

(Transparency)

Another thermoacoustic refrigerator. This one holds a couple of records. It is really easy to set records in a new field like this. This is the current record holder for cooling power. Dr. Garrett is the owner of this refrigerator, at Penn State. It provides 300 watts of cooling power at the freezing point of water. It is essentially a half-wavelength from this loudspeaker through the bent tube back to that loudspeaker.

This is the refrigerator of 300's, easy to remember. It is about 300 cycles a second, 300 watts of cooling power at the two cold heat exchangers on these ends of the two stacks, 300 watts of electric power required to drive it. The gas was 20-bar helium-argon mixture, 20 percent argon, 80 percent helium.

DR. GARDNER: I want to know if this refrigerator also holds the record for the greatest number of physicists attached to it by their tongues.

DR. SWIFT: Are you referring to the famous tongue photograph?

(Laughter)

Those of us who grew up in cold climates know -- we learned the hard way, when we were kids, not to do this. But if you grow up in Los Angeles, I guess you think that it would be nice to taste the frost.

STUDENT: You said 300 watts output and 300 watts input?

DR. SWIFT: Yes. That sounds too efficient to you? It is a good lesson. We should go back to viewgraph one.

(Transparency)

The coefficient performance of a refrigerator is bounded by this ratio of temperatures. Suppose you are at 300°K , which is room temperature, and suppose you are only trying to lift heat by 30° , from the ice point to room temperature -- 300, 30, 10. Carnot says that you could have a COP as high as 10, where the heat pumped would be 10 times the work put in. It does take you aback the first time you see it, but it's true. And so Dr. Garrett's refrigerator is only about one-tenth as efficient as Carnot would permit.

(Transparency)

So those are some examples. That is probably a quarter of the decent examples that exist among attempts at practical thermoacoustic engines and refrigerators. Overall, I would say that the advantages of this technology are that it is going to be reliable and low-cost. There are not any moving parts. That should aid reliability and low cost.

Shortcomings: Relative to the heat engines and refrigerators that are in widespread use in the world, like the steam turbine and the internal combustion engine and the Rankine-cycle refrigerator -- those are what our whole society operates with, basically -- thermoacoustics, as of now, tends to be a little big and a little inefficient.

I want to make a little sermon about the tradeoff between cheap and inefficient. To some extent, you can always make something more efficient by spending more money.

(Transparency)

This is my short version of that lesson. Think of refrigeration technology. Here are a couple of refrigeration technologies, the Stirling refrigerator and the vapor-compression, or Rankine-cycle, refrigerator. Stirling is, in principle, more efficient. You can build Stirling refrigerators that are very near Carnot efficiency. You cannot come close to Carnot with vapor compression. But the vapor compression hardware is so cheap to build. What do we as a society, what does industry choose to do? What is in your kitchen refrigerator? It is a vapor-compression unit. Your air conditioner? A vapor-compression unit. Why? Because in the great tradeoff between capital cost and operating cost, we choose to operate somewhere around here. It is not really so important that this is not the most efficient technology in the world; it is more important that it is low-cost to manufacture.

I think what we are doing in thermoacoustics is this: We are giving the world a new choice of technology. So far, we have not gotten the efficiency quite up to the level of vapor-compression technology. But we are moving this point up as time goes on. Nevertheless, there will be some applications for which people will probably choose to go over to this curve, because of low cost.

MR. WANG: What is the actual scale on the efficiency?

DR. SWIFT: I do not know. It would be a different scale for a kitchen refrigerator and for an air conditioner and a car air conditioner.

MR. WANG: On your graph it looks like high is equal to 80 or 70 percent.

DR. SWIFT: In the large air-conditioning units that are sold for air-conditioning big buildings, like campus buildings at the university or office buildings, Trane Corporation sells some that are about 85 percent of Carnot. They sell only one a year, because they are so expensive that very few people want to buy them. Most of their sales for big buildings are at 75 percent of Carnot. For air conditioning, maybe this is 85 percent of Carnot and this is 75 percent of Carnot and 60 percent of Carnot. By the time you get down to a window air conditioner in your home, it is not so good.

(Transparency)

That was the introduction. Here is an outline of the talk. We will spend quite a bit of time looking at the microscopic behavior within thermal penetration depths and viscous penetration depths, because it is fun and because it is related to a lot of other acoustics stuff, related to a lot of the talks we have heard this week.

Then we will spend some time talking about power and efficiency, because that is where all the action is from the heat engine point of view. I will show you an example in some detail. Then we will talk about what it is that we need to understand that was beyond the understanding of Nicholas Rott in 1969.

(Transparency)

There is going to be a little bit of math here. We will ease into the math by starting with a warm-up problem, something that I think will be intuitively familiar and have easy math: Viscous dissipation and oscillatory flow along a boundary.

Suppose you have a boundary here and there is oscillatory flow in the x direction, and we want to understand the details of that motion. Here is Newton's law for fluids in all its glory: mass times acceleration equals the sum of all the forces on the element of fluid -- pressure force, viscous force. We will have the boundary condition of zero velocity on the surface.

We will use the usual acoustics assumptions. We will let every variable -- pressure, density, velocity -- have some mean value plus an oscillatory part, except the velocity does not have any mean value. Substitute these into here and throw away any terms that are not just subscript 1. Boy, are we lucky; it really gets simple.

Everything goes away except this. The velocity along x can be thought of as a function of y . This is a differential equation for u of y , which you can solve in just a few minutes. This is the solution.

Now it gets hard for us to try to understand it. I just hate looking at these functions of complex variables and trying to remember that really what we mean in acoustics is that the real solution is the real part of this whole thing. I have made a computer animation to show you what that function really looks like in time and in space.

The main point to remember from this viewgraph, when we go to the computer animation, is that the scale length in the y direction is this viscous penetration depth. That is twice the viscosity divided by the frequency and the density.

(Computer simulation)

Here is a resonator full of gas. A solid wall is here. Think of a long pipe, closed at one end, and it has a loudspeaker or a moving piston or something at the other end. We want to look at the half-wavelength standing wave in there. We can just stare at this for a while and enjoy the phasing in a standing wave. This is something that we have all worked with mathematically in all our classes, but sometimes we do not really see it, except when you watch water sloshing in a tub or something.

The pressure is a half-wavelength cosine function in space. Velocity is a half-wavelength sine function in space. There is a little bit of a non-zero velocity right at the end here, because the piston is doing some work to overcome viscous losses in here, which we will get to in a minute.

If you put some markers in the gas -- think of little lines of smoke that would move with the gas -- you would see something like this, where the motion is the greatest in the center, where the velocity is the greatest, and approaches zero toward the ends.

Down at the level of the viscous penetration depth here, at the wall, it must get more complicated. We cannot have non-zero velocity right up against a wall, so we have to look very much closer. The next frame is a magnified view of the yellow zone in this frame.

This shows the solution that I showed you on that last viewgraph, that exponential with complex argument. This is what the motion of the gas really looks like in acoustics, right up against a solid boundary. This moving, wiggly line is one of the smoke lines in the gas. That is how the equation of motion says that it moves.

Each tick mark here is one viscous penetration depth. You have to get out to about four penetration depths before it really gets good and flat, gets into the long-distance limit. This shows the phase and the magnitude and all that stuff from the equation on that viewgraph.

Of course, there is power dissipated in viscous flow.

(Transparency)

Back to the math. Here is that complex solution again for the velocity as a function of y . The instantaneous dissipation of acoustic power into heat in something like this is viscosity times the square of the velocity gradient. If you work that out for this expression, you find that the dissipation per unit volume is largest at the surface where the velocity gradients are the largest, and it falls off cleanly as an exponential as you go away.

At this point, I think you see what the philosophy of this talk is going to be, and, indeed, what the philosophy of most of the talks at the summer school has been. Developing intuition is really important, and we do our best to help you develop intuition, but you have to have math or experiments to back it up. It is hard in a confined time like three hours to really do justice to any of this stuff. The viewgraphs with equations on them are nothing but outlines of calculations that may have taken a graduate student an entire Ph.D. career to work out the steps in between. Things like this have been checked out by experiment, too, such as using laser Doppler velocimetry, looking at the velocity as a function of position away from the wall. You have to take this viewgraph as just an outline of a lot of work that a lot of people have done.

(Transparency)

We are going to move now from the viscous penetration depth to something closer to thermoacoustics, the thermal penetration-depth problem. Here is a simple thermal penetration-depth problem. Suppose we have a solid wall and we have oscillatory motion toward and away from the wall, so there is basically an oscillating pressure at the face of this wall.

The wavelength is huge compared to this whole picture. We only care about what happens within about a thermal penetration depth of the surface. If you write down the equation of heat transfer for a gas, in all its glory, it is all this stuff. But again, the acoustic approximation -- the usual substitutions that we always do -- makes it really simple.

That looks a lot like the differential equation that we had for u_1 as a function of y . Here, we have T_1 as a function of x . It is the same differential equation. It has the same kind of solution. The temperature, the oscillating temperature, near this boundary has this complex solution, with the characteristic length being the thermal penetration depth, the square root of 2 times thermal conductivity over frequency and heat capacity per unit volume.

So back to the animation.

(Computer simulation)

Our standing wave again, the same as before, gas moving back and forth. Now we want to look in detail at what goes on here, where the pressure is oscillating, the gas is moving straight toward and away from the wall just a little bit. We know that this picture cannot be correct, that the temperature cannot be going up and down right at the wall, because if the wall has the heat capacity of typical solid walls, it is going to thermally anchor the oscillating temperature right at the wall.

Out here in the middle of the resonator, much farther than a thermal penetration depth from any boundary, the temperature oscillations follow the pressure oscillations. But that cannot be true right near the wall. The next frame shows what happens right near the wall.

It is the same solution as for the viscous flow problem in the last animation. It is turned on its side now. Temperature here, as a function of position. Each tick mark is one thermal penetration depth. By the time you get about four penetration depths away from the wall, you really do just have adiabatic temperature oscillations, with the temperature going up and down in phase with the pressure. But right at the wall, the isothermal boundary condition imposed by the wall anchors this at zero. The differential equation tells us that is what happens in between.

With the viscous problem, we looked at power dissipation. It is intuitively pretty clear that in viscous flow you should get power dissipation, because it looks like friction. It looks like fluid layers sliding relative to one another. It is not so obvious where the power dissipation comes from here in the thermal penetration depth problem.

To try to illustrate that, I want you to think of an imaginary piston that moves with the gas at some point and ask whether or not that piston does work on what is in front of it. If it does work on what is in front of it, then there must be some source of dissipation there. If you want to calculate the work that a piston does, you calculate the integral of $p \, dV$. The horizontal line is p and the front of the piston marks V , so the intersection of those two lines traces out a little reciprocating line -- no, not a line, it is a little loop!

The area of that loop is the time integral of $p \, dV$ for that piston. The piston is, in fact, doing work on the gas in front of it, so there must be some source of power dissipation in there.

The origin of that power dissipation is the heat flow that takes place along the temperature gradients in there. Suppose I pause this right about there -- the gas right up against the wall here is not experiencing any temperature oscillations. Its temperature is fixed. Pressure times volume for a little element of gas there is constant -- no big deal. Gas out here is adiabatic. That is the ordinary acoustics we always deal with. Not a lot of dissipation there.

But gas right about in here is really confused about whether it should be adiabatic or isothermal. It feels stiff to the piston when it is getting compressed and its temperature is going up, but then some heat runs down the temperature gradient to the wall, and it cools and shrinks. Later in the cycle, the opposite occurs. The heat comes back out of the wall and the gas heats and expands. So that is the source of the dissipation here: thermal expansion and contraction, suitably phased with the pressure.

If you make a plot of the area of this loop as a function of where you put the imaginary piston, this is the way it looks. Right in here, where the slope is greatest, that represents the most dissipative gas. That is the stuff where the given dx of gas is making the biggest change in work.

MR. POESE: Why is the slope of the p dV curve the biggest at the point of most dissipation?

DR. SWIFT: Because this p dV is the work done on all the gas between the piston and the wall.

MR. POESE: I see. So after four thermal penetration depths, there is no more gradient, so there cannot be any more heat, so that is why it is constant?

DR. SWIFT: Right. Then there is no more additional dissipation. Good question.

(Transparency)

If you look at this from the point of view of one little parcel of gas -- here is a little parcel of gas that is shaded red when it is at its right-most extreme of motion and shaded blue when it is over to the left and it has been compressed to a smaller volume by the sound wave -- if that little parcel of gas is really close to the plate, if this distance between here and here is much less than a thermal penetration depth, then the pressure and volume as a function of time for that parcel of gas follow an isotherm, $pV = nRT$; pV equals a constant, back and forth. There is no area to that path, so there is no dissipation coming from that element of gas.

In this case, if this distance between here and here is much greater than a thermal penetration depth, that gas is insulated from the wall. It does not feel the presence of the wall at all. That is the ordinary sound case, adiabatic sound. This line is a small piece of p V to the gamma equals a constant, an adiabat.

But if you are in between, where the parcel of gas you are looking at is about a thermal penetration depth away, then you get this kind of pressure-volume diagram for that gas, non-zero area, dissipation.

DR. HOFER: I am confused. If we go back to the simulation, it seems that the fundamental dissipation is not really the per-unit volume of gas. Isn't it really the magnitude of the time-averaged temperature gradient? If you look at the simulation, the temperature gradient is highest right at the wall. In that other curve you showed, was that the trace of the piston as you moved the piston forward?

DR. SWIFT: Is the first question, isn't the dissipation really proportional to the square of the temperature gradient? So why does this appear to come in with reduced slope here?

DR. HOFER: If you look at the top curve, it seems clear that the highest temperature gradients are right at the wall.

DR. SWIFT: That is right.

DR. HOFER: The bottom curve is what, again?

DR. SWIFT: The value at this point is the area of this pV loop that this piston experiences here. If we had put the piston on this line instead, we would get a little bit smaller pV work. This curve shows how much total dissipation occurs to the left of the given x .

DR. HOFER: But as you move the piston in, you are reducing the total amount of volume that is being compressed, so that is really non-indicative of the per-unit volume.

DR. SWIFT: Right. It's the slope that is indicative of the dissipation per unit volume. Out here, the slope is zero, which indicates that between this point and this point you are not picking up any additional dissipation. Right here, where the slope is the greatest, that is the place of maximum dissipation.

DR. HOFER: If the dissipation per unit volume were constant, then you would get a straight line with a constant flow, wouldn't you?

DR. SWIFT: Yes.

DR. HOFER: So you really have to factor that out if you are looking for differences in the per-unit volume dissipation. Just looking at the top curve, it seems that the most dissipation per unit volume is where the temperature gradient is the highest, which is at the wall.

DR. SWIFT: Why do you say that? Why is the most dissipation per unit volume where the temperature gradient is the highest?

DR. HOFER: Because that is the source of irreversibility, heat flow through a diffusive medium.

DR. SWIFT: I think this is a really subtle question. It has to do with to what extent in a given unit volume the entropy generation in that unit volume, which is proportional to the square of the temperature gradient, is equivalent in that same unit volume to the work dissipation, which is the cycle integral of $p \, dV$. That is a tricky question.

There is a really fun book that just came out called *Entropy Generation in Heat and Fluid Flow* by Bejan. A third of the book has to do with stuff like this.

DR. KEOLIAN: Where the temperature gradient is the highest, the heat flow will be the greatest, but the -- I want to say that it is isothermal there and so there is not any -- DR. SWIFT: This is the sort of thing we could argue about as long as you argue about Helmholtz resonators, I think. It is really tricky. Do the math.

MR. WORLIKAR: (Inaudible.)

DR. SWIFT: Yes, exactly a standing wave and no phase difference. It is very nearly that. This is very nearly just a reciprocating straight line. It is a subtle effect.

(Transparency)

We can connect all this to things we have seen elsewhere this week or in our textbooks. If you look in Kinsler and Frey's acoustics text, the bulk attenuation of sound is written this way. You can rewrite that in terms of our viscous and thermal penetration depths in a way that I think is sort of pretty. It ends up showing the ratio between the penetration depth and the wavelength everywhere.

The reason that ordinary sound propagation in air is not very attenuated is that the distance that momentum or heat can diffuse in a quarter cycle of the sound wave is a lot smaller than the wavelength, which is the distance over which temperature variations and so on exist.

You will also find attenuation in pipes, where the viscous and thermal contact to side walls in a propagating wave is more important than bulk attenuation, expressed like this. You can take the ratio of these two expressions if you want to find out under what circumstances side wall damping is more important than bulk damping.

(Transparency)

Finally, we are done with warm-up problems. Finally, we are going to get on to thermoacoustics. This is the simplest thermoacoustics example that I can think of. It ignores viscosity. We have never come up with any good intuitive picture that includes viscosity. But we do know how to do the math with viscosity, so do not be alarmed at the lack of viscosity in our intuitive pictures.

Here is a zero viscosity, simple thermoacoustics problem, back to oscillatory motion along a wall. This time we will have a nonzero temperature gradient along the wall. Our linearized equation of heat transfer again -- it has a new term in it relative to the one we saw a couple of viewgraphs ago, having to do with the velocity and the temperature gradient. The temperature gradient was zero the last time.

You solve this for T_1 as a function of y , and you get this mess, the same complex function of y that we have been looking at all along, and two terms in the magnitude. One of these terms comes from adiabatic oscillations. If the pressure is oscillating in a gas, then the temperature is oscillating, if there is no thermal contact. That should be familiar to you.

This term you may have never seen before, but it is really easy to understand. If you have gas with a temperature gradient moving along x -- suppose the stick here is gas, with temperature being a vertical axis and position being a horizontal axis, and the gas is moving, carrying the temperature gradient with it, and you sit at a fixed x and look at what happens, you see the temperature go down and up, and down and up. That is all this term is.

So those two sources of temperature oscillation, moving gas with a temperature gradient and oscillating pressure, combine with this complex function that accounts for thermal contact with the walls, to make such a monstrosity that you get extremely rich behavior.

If you go to big distances from the wall, you get one thing. If you go really close to the wall, you get something else. If you go in-between, you get something else. Depending on the relative size of this term and this term, you can get positive, negative. Depending on standing-wave phasing versus traveling-wave phasing, for those of you in the back row, you get added variety here. All of those things are fair game in thermoacoustics.

(Computer simulation)

We will do one more animation to capture some of that. Once again, a standing wave in a pipe, half-wavelength. Now put in a set of solid plates lined up with the acoustic oscillation.

We did an animation where we thought about what happens here, where there is velocity but no pressure, so that the only thing that really matters is the viscous shear and the momentum equation. We did an animation where we looked over here, where the only thing that really mattered was the equation of state and the continuity equation. Now we are going to be somewhere in between, where it is going to get really complicated. We had better do a magnified, close-up view of what happens right in there.

Here is the close-up view. This is one little plate in the stack, between there and there, and here is another little plate in the stack between there and there. We will keep our eye on one parcel of gas. Here it is moving back and forth up there, shown magnified down here. You recall that when the gas moved that way, the pressure was going up. So it must be that the density was going up, so the parcel should be smaller when it is here, bigger when it is there. Now we start thinking about pressure-volume diagrams and that sort of thing.

If that gas is one or two thermal penetration depths away from the solid boundary, then it is experiencing the same kind of oscillatory thermal relaxation as in the last animation. There will be a non-zero line integral of $p \, dV$. That parcel of gas is absorbing work from the sound wave.

You can understand the source of that absorption of work by thinking about heat flow between the parcel of gas and the solid, and temperatures. Here, the horizontal blue line is the temperature of the solid, and the white trace shows the temperature of this parcel of gas as a function of time.

The standing wave moves the gas left and pressurizes it, and so it heats almost adiabatically. The temperature rises above the temperature of the solid plate. But it is not really far from the plate. A little bit of heat can then diffuse to the plate -- that is the red arrow -- cooling the gas. That is why the volume drops a little bit and opens up the pV oval a little bit at this end.

Then the standing wave carries the gas back the other way, cooling it adiabatically, making it colder than the plates, allowing thermal expansion at that end.

You really have to stare at this a long time and try to keep five things in your mind at once. It is not easy. But this picture explains a lot. It explains the direction of the heat flow at different phases of the cycle, and it explains the fact that the gas shrinks thermally at high pressure and expands thermally at low pressure. It gives you the sign of the integral of $p \, dV$.

From the point of view of work, this gas parcel is absorbing work from the standing wave. From the point of view of heat, this parcel of gas picks up a little bit of heat from the plate here and deposits it over there at a slightly different position. All the little parcels of gas in here are doing that. They are all absorbing work, and they are all moving heat from right to left, each cycle

of the sound wave. So now we are starting to get something that looks like a heat pump or refrigerator.

The only thing I have to add to make a heat pump or a refrigerator is a couple of heat exchangers, one on either end of the stack, so that we can communicate this heat transport to the outside world.

If you allow the temperature gradient in the plate to be non-zero, but still small, you do not change this picture qualitatively at all. The gas is still cooler than the plate at the right extreme of its motion and hotter than the plate at the left extreme of its motion. The direction of heat flow is just the same as it was before -- a little bit smaller, because these temperature differences are smaller, but not a big deal.

The area of the loop is a little bit smaller because the thermal expansions and contractions are a little bit smaller, because the heat flow is a little bit smaller. Qualitatively, it is the same picture as before. But now we have a whole refrigerator.

The parcels of gas absorb work from the standing wave. Each parcel of gas moves heat a little bit to the left. The first parcel picks up some heat from here and moves it to here. The next parcel picks it up from there, moves it to there. The next parcel picks it up right there and moves it to there, and so on.

MR. HALLAJ: Are the plates in the center going to experience this more than the ones on the side? Near the center of the stack you are going to have two people putting in heat, say, from above and below the plate you are looking at -- or more, obviously, if you are in the middle. It is like a conduction problem.

If you drew another parcel of gas on the other side of one of these plates, you will see that there are two parcels of gas depositing and taking heat from one of these plates. In heat conduction, it is analogous, probably, where a few more plates up, you have lots of plates depositing and taking out heat, lots of parcels.

DR. SWIFT: Each one just has two sides.

MR. HALLAJ: The plates are identical?

DR. SWIFT: Yes, except in this direction, of course, they are not identical. If you have a square resonator, they are all identical.

DR. GARRETT: The easy way to think about it is to consider your plates there as just being half-plates. You are showing only half the thickness of the real plate, and everything is symmetrical.

DR. SWIFT: Dr. Garrett is saying to think of this line as a symmetry line in the middle of the plate and then flip the picture over across that mirror-image plane.

MR. LANDSBERGER: I think we heard before that a particle that goes back and forth can have about a 6-degree temperature difference. Is that about right?

DR. SWIFT: Usually, these thermoacoustic devices have oscillatory pressures that are 5 or 10 or 20 percent of mean pressure, so let us say 10 percent. Then the oscillating temperature is 10 percent of the absolute temperature, with some factor of $\gamma - 1$ or γ that I can never remember, which reduces it a little bit. If you are at 10 percent pressure oscillations and at 300° K ambient, you have 30° K temperature oscillations.

MR. LANDSBERGER: Okay, that is the extreme front end of a heat exchanger. But you are saying that one right there --

DR. SWIFT: No, no. That is the difference between here and here.

MR. LANDSBERGER: At the two ends of the stack, then, it should be higher?

DR. SWIFT: Right. I think the temperature record for a refrigerator is held by one of Tom Hofler's refrigerators, dry-ice temperature or colder than that.

We started with the temperature gradient in the plate zero, and now we have gone to a non-zero temperature gradient like this. What is going to happen if we let this thing develop a temperature gradient about like that? If the temperature gradient is a little bit steeper, then the temperature gradient in the solid exactly matches the temperature gradient that the adiabatically moving gas experiences, because when it moves to the left the pressure is going up, so the temperature is going up, and when it moves to the right the pressure is going down, so the temperature is going down. So when the temperature gradient has this value, no more heat transfer, no more thermal expansion or contraction to open up a pV loop, no nothing.

What happens if we let the temperature gradient be steeper? Minus sign on all previous pictures! Now the temperature gradient in the solid is steeper than the average slope of the oval, steeper than the gas motion with adiabatic temperature oscillation. Now, when the gas goes over to the left, high-pressure extreme of its motion, even though it has been adiabatically heated and its temperature has gone up a little bit, it is not hot enough to match the local plate temperature. The plate is hotter than the gas. Heat flows from the plate to the gas. Thermal expansion makes the volume increase at that time.

Now the sign of circulation on our pV diagram is the opposite of what it was before. So the line integral of $p \, dV$ is positive. That parcel of gas is doing work on its surroundings, pumping work into the standing wave. In the case of this demonstration, that is what was happening. That work going into the standing wave was then being radiated from the standing wave out the end of the tube to fill the room with sound.

The bucket brigade of heat now is picking heat up off the plate at the left, depositing it on the right. That is also the opposite of what was happening before. Now heat is being carried by the gas from hot to cold.

Before we take a break, I want to go through a few more viewgraphs to reinforce some of the math behind this.

(Transparency)

What do we do in acoustics? We start with fluid mechanics and then we get simple. Full-blown Newton's law, within the acoustic approximation, turns into a simplified equation like this that we can treat as a differential equation for the x component of the velocity as a function of dimensions in the perpendicular direction.

You can solve that differential equation in different geometries -- circular pores, parallel-plate geometry, triangles, whatever you want. We are not going to do that. I have not done much of that myself, but other people have done that, and you can look those solutions up. There are exponentials with complex arguments, trigonometric functions with complex arguments, Bessel functions with complex arguments, series solutions with complex arguments.

You can integrate those solutions for u with respect to y and z , to get a spatial average of velocity, and rearrange the result to make it look like Ohm's law. So Newton's law turns into Ohm's law, in my mind, in thermoacoustics: Here is a pressure difference across some length of stack or duct or something, being proportional to the current, with a complex proportionality constant that is like a complex impedance in Ohm's law. This has inertial parts and resistive parts. They are like inductive and resistive in electrical circuits.

All the details of this differential equation and the boundary conditions and everything else get wrapped up in something that has come to be called f_{nu} . Other people have worked out what f_{nu} is, and you can just look it up.

(Transparency)

Here is what f_{nu} is for some different kinds of geometry -- parallel plates, circular pores, rectangular pores. I think I chose a 6-to-1 aspect ratio here just to be specific. This is the boundary layer limit that we have been working with on most of the viewgraphs, the e to the minus 1 plus $i y$ over δ . We will see this again in a little bit. One thing to notice is that it does not really matter a whole lot what kind of geometry you work in.

(Transparency)

Another starting equation for acoustics is the continuity equation, or mass conservation equation. If you do spatially average it over a pore in a stack, you get a simpler version. If you include with that the ideal gas equation of state, which says that the density depends on pressure and temperature, and watch your spatial averages -- the spatial average of the density, the spatial average of the temperature -- you can do a little bit of algebra and get this version of the continuity equation, which says, to me, that the reason that the flow in thermoacoustics is not like hydraulic fluid, with a slug of fluid moving past our observation point with the same velocity going out of the box as came into the box -- the reason it is not hydraulic fluid is compressibility and thermal expansion coefficient. Those are the realities of the gases we deal with.

Here is the time dependence. Here is the compressibility of an ideal gas. The f_k has to do with whether we are dealing with an adiabatic compressibility or an isothermal compressibility, or something in between. This f function is the same as the viscous f function we just talked about a minute ago, but with thermal penetration depth instead of viscous penetration depth.

Compressibility term, thermal expansion term: If you have moving gas in a temperature gradient with thermal contact to the walls, you are going to have thermal expansion and contraction. It is very complex and detailed. But it leads to a different velocity coming out of your box than what was going into your box.

(Transparency)

For historical reasons, we always like to put up Rott's wave equation. That is what results when you take these two equations we just looked at and eliminate u . That is what we do in wave equations, eliminate u between the momentum and the continuity equations, to get a second-order differential equation for pressure. I find it easier to look at these two first-order equations separately than looking at Rott's wave equation. It just seems a little bit more intuitive to me.

For the next part of the talk, when you look at these thermoacoustic engines and refrigerators, I want you to imagine the momentum equation and the continuity equation in action.

I want you to imagine, in the resonator, momentum carrying gas back and forth, and the continuity equation and the equation of state providing the spring that keeps the resonance going.

Within the stack, and everywhere else, I want you to imagine these thermal and viscous penetration depths, where all the really intricate, interesting action is going on. These are very tiny -- submillimeter, typically. Momentum and heat are diffusing laterally, back and forth.

Try to imagine all that, and then try to imagine that you are out of graduate school and you have your Ph.D. and you have a job, and you are working at Ford or Whirlpool and you are building a thermoacoustic refrigerator. Try to imagine what your boss is going to be like at a place like that.

Do you imagine your boss is going to care a whole lot about your detailed understanding of viscous and thermal penetration depth activities? No. The boss cares about final result -- what is it going to cost to build it, what is it going to cost to run it, can you do it yesterday, how much power can you cram into a package this big, and how efficient is it going to be?

We have not talked about power and efficiency at all yet, which are the things that are important from the practical or business point of view. Power and efficiency are next.

(Transparency)

We start with acoustic power, something you are all familiar with. We know that the acoustic power is a second-order quantity equal to the area of the pipe that we are dealing with times the time average of the oscillating pressure and the spatial average oscillating velocity.

If we are trying to ask how much power gets absorbed (or produced, in the case of engines) in a length dx , we can write down dW_2/dx and use one of the most important techniques in mathematics to take this derivative and break it up into two terms. The d by dx of pu is dp/dx times u plus p du/dx .

From half-an-hour ago, we know what dp/dx is, from the momentum equation. That was the one that ended up looking like Ohm's law. We know what du/dx looks like from the continuity equation. We can stick those expressions in and do a half-hour's worth of algebra and get this mess for the work absorbed in length dx of a duct or a channel within a stack, or something like that.

Let us look at these terms for a second, just to give them an overview. Here is a term that is proportional to velocity squared and it has f_{nu} in it. So this has to be a viscous term. The dissipation you get in oscillatory viscous flow is proportional to velocity squared -- all right, this is the viscous dissipation term. It always absorbs power, because the sign is always positive, because you have the square of things and you have the imaginary part of $-f_{nu}$ and it is positive.

This one has f_k , so it is a thermal relaxation term. This one has a little bit of both, but it definitely needs some thermal relaxation to be nonzero. These are the thermal terms. Again, you can see that the important thing here is the imaginary part of f_k . It shows up obviously here. Here it looks like the real part, but wait a minute.

In a standing wave, u and p are 90 degrees out of phase from each other. So the important contribution to the real part of something times something that is imaginary is the imaginary part there.

So here is where all the action is in thermoacoustics. It is in the details the imaginary part of minus f_k . Basically, it turns out that you want to make this as big as possible. You want to work somewhere down at the bottom of that curve.

Viscous effects are always bad. We need the thermal effects to make it work. How can we improve our efficiency? It would be nice to be able to make this term smaller compared to the others. Because of the Prandtl number of real gases being pretty close to 1, f_{nu} is always pretty close to f_k . So you can do a little bit there by changing gases, but not very much.

Another thing you can do is try to make u_1 small relative to p_1 . You do that by moving the stack in the standing wave to a place where velocity is a little smaller and pressure is a little bit higher. You see that there.

The final point from this viewgraph is that you can think of the work flux as a function of x in one of these thermoacoustic devices as a curve like this, and dW/dx , the slope, the steepest in the places where all the interesting action is, in the stack or whatever.

That is acoustic power, but there is more to power in thermoacoustics, because we have heat as well as acoustic power.

(Transparency)

Here are some things that we have not encountered in this summer school yet. The question is, what do you really mean by total energy, or energy flux, in a thermoacoustic situation? The one thing that we are pretty confident of as physicists is energy conservation.

If we draw a control volume, an imaginary boundary, here and invoke the first law of thermodynamics at that boundary and say what goes in must come out, the work that goes in must equal the sum of the heat out and the energy out. Well, what do we mean by energy flow down through the stack like that? We will come back to that in a second.

(Transparency)

Here is another control volume. Energy into that control volume must equal energy out of that control volume in steady state. It is not going to pile up in between. So this tells us that, whatever definition of energy we come up with in a minute, it had better turn out that the energy flux within the stack is independent of x . After all, the side walls are insulated, so heat cannot go in or out, and the side walls are not moving, so work cannot go in or out. What goes in must come out; energy flux has to be constant.

(Transparency)

What do we mean by energy? In fluid mechanics, the enthalpy is the correct energy to use. There are derivations of this in any do good thermo-hydro book, like Landau and Lifshitz. It is the correct energy to use because it includes both the internal energy and the pressure. The pressure is part of the energy flux, because when a parcel of fluid is moving past some place, it does work on the fluid in front of it.

If you take the energy flux density here, kinetic plus enthalpy, and do the usual acoustics approximation, keeping things only to second order, you get about nine terms, but all of them go away except this one, the time average of the product of oscillating enthalpy and velocity.

There are two ways to break that down using thermodynamic relations. Neither one is easy, but that is life. If you express enthalpy in terms of entropy and pressure, and stick that in here, you get an energy flux density that has something that looks like entropy transport -- maybe, loosely, we could call that heat transport -- and something that looks like our old familiar acoustic intensity. So it is appealing to talk about energy flux in thermoacoustics as being the sum of a heat flux and a work flux.

Another way to look at it is shown here on the right, where we express enthalpy in terms of temperature and pressure. For an ideal gas, the coefficient in front of the pressure term turns out to be zero. This one is very simple to work with computationally. The energy flux density is just heat capacity time the time average of temperature and velocity. Nicholas Rott used that expression to write down an expression for the energy flux in thermoacoustics.

(Transparency)

We are not going to look at the details of this too much, but I want to show it to you so that you can see that it is not very threatening. There are a few terms that we can recognize right away, like this 1 right here times p_1 and u_1 . That is the ordinary acoustic intensity. Here is ordinary thermal conduction in the x direction. That is not very difficult or interesting.

Then there is this term that is velocity squared times temperature gradient. That is energy flux due to the bucket-brigade picture that we had, with gas moving back and forth along the temperature gradient picking up heat in one place, putting it down in another place.

The remaining term is a little bit harder to interpret.

Think about that equation for energy flux, total energy flux, as a differential equation for the mean temperature as a function of position in the stack. If the stack is thermally insulated on the side walls, energy into any control volume of length dx has to equal energy out. Thus H_2 has to be constant, as we said before, within the stack. It is some number. So we have a really complicated differential equation for T_m of x . T_m appears explicitly/ here and here. There are T_m 's everywhere in here, because viscosity and thermal conductivity and density and everything else depends on T_m . It is a mess, but it is not too bad in a computer.

(Transparency)

This brings us to the philosophy of some of the most widely used thermoacoustic calculation methods. From the first segment of the talk, we had the momentum equation and the continuity equation, the "Ohm's law" and compressibility and thermal expansion pieces of things. We can let somebody else work out what f_{nu} and f_k are once and for all, and get all the Bessel functions right. Then we can do a numerical integration of these equations as a function of x .

We do not have to worry about the y and z dependences anymore. We use the equations from the first part of the talk for dp_1/dx and du_1/dx . We use the energy equation that I just showed you for dT_m/dx . Now we can do a numerical integration along x and find out all kinds of interesting things.

Software to do is documented and available. You can get it on our Web site. You can also get the animations that I showed you an hour ago on our Web site if you want them.

There are some approximate calculation methods as well that usually start with something like this and make different simplifying approximations. Some of those are easy enough to do with a spreadsheet or a pocket calculator or even on a blackboard if you are really brave.

DR. MIGLIORI: Greg, on the boundary-layer approximations for f_k and f_{nu} , how good is that? You must have checked this.

DR. SWIFT: Do you remember the graph of f versus pore size over penetration depth? -- that tells how good it is.

DR. MIGLIORI: Rather than using the real solution between two plates?

DR. SWIFT: It is not too bad. By the time you get to realistic conditions, where the spacing between plates is three or four thermal penetration depths, you have it to 10 percent, probably.

DR. GARRETT: The approximation techniques known today work extraordinarily well, but they do not include a traveling wave part. A typical approximation would give you only assumed standing waves. If you do take that extra step to letting a computer solve it for you, you pick up that bonus from the traveling wave effects. But as you point out, it is surprisingly good. The simple approximate boundary layer stuff is extraordinarily good.

DR. SWIFT: Probably none of us starts with the full-blown computer calculation. You always have to work with simpler math to get oriented overall and get your feet wet.

STUDENT: Dr. Swift, what is the Web address?

DR. SWIFT: It is <http://rott.esa.lanl.gov/>.

I want to take a little bit of time now to show you a specific example of real hardware that we designed and built, and how it all worked. For that, I will go back to this example of the thermoacoustic engine we built for Tektronix Corporation to drive their pulse-tube refrigerator.

(Transparency)

I showed you this before. It is a half-wavelength engine along here, driving this pulse-tube refrigerator, which we really do not have time to talk about too much. Another aside for the people in the back row: You might remember somebody who is missing from this picture at Tektronix. Two years ago at the summer school, Charles Jin was a discussion leader. What happened to Charles Jin? Before this picture was taken, he quit Tektronix and went to work on Wall Street. He is a mortgage banker now.

(Transparency)

Tektronix approached us. They wanted to make a cryogenic refrigerator with no moving parts for cooling electronics. They said, "Give us the thermoacoustic engine. We will make the pulse-tube refrigerator." They said, "This is what we want. We want it to use 30- bar helium gas at 400 cycles/second. We need a kilowatt of pV power at our pulse-tube refrigerator. We want that delivered to us with an amplitude of 3 bar. Our pulse-tube refrigerator is a complex acoustic impedance. It has about a -30-degree phase. It is mostly a real impedance, but there is some compliance in there. We want it as small as possible and as efficient as possible."

We went back to them and said, "Wait a minute. If we make it as small as possible, it won't be efficient; if we make it as efficient as possible, it won't be small." It turned out that what they really meant was not to make it any bigger than about a cubic foot, and given that, to make it as efficient as it could be.

That was the basic idea. The total envelope of the thermoacoustic engine was constrained by a volume.

So that was the job. We made some seat-of-the-pants decisions -- for example, that it ought to be a half-wavelength long rather than a full wavelength or a quarter wavelength or some of the other possibilities, and that it ought to have an engine in both ends of the half-wavelength. We put that topology into the computer and let the computer tell us how to get the highest efficiency, what stack length, gap -- all that other stuff -- we should choose.

(Transparency)

I'll show you a few things just to give you an idea of what information we can get from these calculations. This is a scale drawing of the hardware, except the details in here are not to scale; but the resonator shape is to scale. Superimposed on that is a plot of the magnitude of the oscillating pressure as a function of position and the magnitude of the oscillating velocity as a function of position.

If you look in the handouts, you really can see a few little kinks occasionally on this curve, but basically it looks a lot like a cosine function. This would look a lot like a sine of x , except that where there are area changes, like in this cone, and at the little abrupt area changes that the gas experiences as it goes in and out of the heat exchangers, there are little discontinuities in velocity when you get to those area changes. Volumetric velocity, we say in acoustics, is conserved at the junction between two acoustic elements. So, basically, it is a standing wave. That is standing-wave physics.

The power issues all come up in the out-of-phase components of pressure and velocity. Now I have relabeled the magnitude of the oscillating pressure as the real part of the pressure. I have multiplied the imaginary part of the pressure by 20 to make it show up on this same scale. The same with the velocity: the out-of-phase component is multiplied by 20. This really is nearly a standing wave. Nowhere in it is the phase between pressure and velocity different from 90 degrees by any more than 4 degrees or something like that.

A lot of times it is good to think of these thermoacoustic engines as just standing waves, if all you are interested in is the overall wave shape. But if you are interested in the power, you have to deal with these out-of-phase parts.

(Transparency)

Some other things you can learn in a design like this include what the temperature is as a function of position. It looks like that in the calculation. It is appreciably different from a straight line. A lot of the approximate calculation methods assume a straight-line temperature profile. That is probably fine for refrigerators which span a small temperature difference, but we are spanning a very big temperature difference here. It is not a great approximation here.

Here are some energy-flow plots. The discontinuity here in total energy is the heat that we supply at the hot heat exchanger, and the discontinuity here is the heat we remove at the water-cooled heat exchanger here. You can see dw/dx here in the stack. At every dx length, the stack is

doing work on its surroundings. Everywhere but in the stack the slope of w is negative, indicating absorption of work.

Our other engine, down to the right, is supplying 500 watts of pV work. This one is supplying 600 watts of pV work. The 500 watts that comes in here is down to 400 watts by the time you get here because of the losses along the resonator wall here. So into the load, where we are supposed to be delivering our 1000 watts to Tektronix's refrigerator, we have 600 watts from the left and 400 watts from the right

-- 1000 watts altogether.

(Transparency)

Calculated performance (we will get to the measurements in a second): To deliver our 1000 watts, we ended up needing 5380 watts of heater power. This includes the thermoacoustics stuff, plus some easily calculated additions, like a couple of hundred watts of heat leaked to the room through the insulation and stuff like that. That efficiency is 18 percent.

If you take the ratio of those two numbers, that is about 30 percent of Carnot efficiency. Why isn't it Carnot efficiency? Why can't it be Carnot efficiency? Our lives would be so much easier if it were Carnot efficiency.

If you make an attempt at an honest breakdown of where the irreversibilities come from, this is how it breaks down. About half of it is thermal relaxation loss in the stack. The very action that makes all this possible, that makes thermoacoustic work generation and heat transport possible, is heat transfer across that thermal penetration depth. It is not isothermal heat transfer. It is lossy; it is irreversible. We are stuck with that.

DR. MIGLIORI: If you had an extremely anisotropic thermal conductivity, what happens? Do you win there in the stack -- thermal relaxation?

DR. SWIFT: You get rid of that one (conduction along x). This (thermal relaxation) is all in the y direction.

DR. MIGLIORI: That is the y direction? Holy cow. I have forgotten how this all works. What did you want to do to make it better? Are you going to talk about that later?

DR. SWIFT: I wish I knew how to make it better.

DR. MIGLIORI: If you could design any material you like, what would you change?

DR. SWIFT: You cannot win that much by changing materials. You just cannot. If I had a gas with zero Prandtl number, I would love that -- no viscosity. Then I could get rid of this.

DR. MIGLIORI: But going up in specific heat or thermal conductivity or something like that does not help.

DR. SWIFT: No, it really does not.

MR. POESE: Isn't viscosity the key to the whole process? Doesn't that help cause the phase difference between the temperature oscillation and the pressure oscillation? If you had no viscosity, would this work?

DR. SWIFT: Yes, I think so. That final animation did not have any viscosity in it. The blob was moving back and forth like a blob. Viscous shear was not changing its shape.

DR. GARRETT: Greg, you had built a prime mover like this using either sodium or sodium potassium eutectic. The Prandtl number was -- yes, terrible. It was so terrible that it was shut down at Los Alamos, where "terrible" has a completely different scale.

(Laughter)

Even there, it was too terrible. The Prandtl number there was a percent or something, right?

DR. SWIFT: The Prandtl number was 0.00-something.

DR. GARRETT: Yes, small, less than a percent. Viscous disappears almost entirely.

DR. SWIFT: It is the square root of the Prandtl number, so it is the square root of "almost entirely." The Prandtl number is low not because the viscosity is low, but because thermal conductivity is so high. That made this (conduction along x) in the fluid itself pretty bad. It was not a good tradeoff.

DR. GARRETT: It was not a lot better than 18 percent.

DR. SWIFT: Plus Tektronix would not want to put it in its oscilloscope.

DR. MIGLIORI: Steve, I still have the piece of Lexan bullet-proof plastic that the sodium engine launched one of its pressure transducers almost completely through.

(Laughter)

DR. SWIFT: The fact that Tektronix constrained us so hard in size is one of the things that made this (thermal relaxation in stack) so big. In order to get a lot of power per unit volume of stack, you have to have a steep temperature gradient relative to the temperature gradient that the adiabatic gas motion follows. That makes for larger temperature differences and more irreversibility. We could do better if we were not constrained by size. But you are always constrained by size.

(Transparency)

I want to show you what the hardware really looked like. I have photographs following this, but this is to get you oriented. This is a scale drawing of one of the business ends. Everything is to scale except the spacing in the stack, which actually would have about 200 lines in it.

This is the water-cooled heat exchanger. It is copper fins with a water jacket. This is the stack. The heat source was an electric resistance heater, a Nichrome ribbon wound on a ceramic frame. The inner pressure vessel here contains the oscillating pressure. It is red-hot, or orange-

hot; metals are pretty weak at those temperatures. We did not want to have to support the whole 30-bar pressure with a red-hot pressure vessel and have to deal with the thermal conduction down along such a thick vessel to room temperature.

So we put this outer shell on, full of thermal insulation and pressurized with argon to 30 bar. There is no mean pressure difference across the inner shell. It just has to support the plus and minus 3-bar oscillations. Our insulation is all packed into this space.

(Transparency)

Here is what the hardware looked like. Here is that inner pressure vessel, with a flange on the end. There are some thermocouples here that have been clamped into position on the outside of the case. We are packing some of the ceramic fuzz insulation around it.

(Transparency)

That ceramic fuzz all got packed in there and then the outer pressure vessel got put on -- here is a spare outer pressure vessel -- and welded shut. We never have to deal with insulation again. It is inside a welded-shut container. Here is the hole where we are going to put the stacks and heat exchangers and so on.

(Transparency)

First, the hot-heat exchanger: aluminum ceramic parts, holding a Nichrome ribbon that weaves back and forth, back and forth, back and forth. That is all held together with superglue so that you can handle it. But once it is in there and the stack is put in on top of it and you heat it all up, the superglue just turns into methane or something that comes out. But the heat exchanger is trapped, so it does not fall apart. Then, if you have to take it apart, you take the stack out and all these pieces fall out.

(Transparency)

We had two stacks. I will show you data from both of them in a second. This is a honeycomb. The pore size is some fraction of a millimeter. It is not the greatest stack, but it is commercially available, readily and cheaply. These three, all stacked in series, formed what we called the honeycomb stack. They are made of Hastelloy, which is a nickel alloy that looks and feels a lot like stainless steel, but it is not quite as soft as stainless steel at the temperatures we were at.

(Transparency)

Then we had the real stack, which was a parallel-plate unit wound as a spiral. This is 3 inches long, 2-1/2 inches in diameter. The Hastelloy sheet is 2/1000ths of an inch, or 50 microns, thick. The gaps are 10/1000ths, or 0.25 mm.

The way we made this was to co-wind the Hastelloy with sacrificial copper, to make a jelly roll, and then cut little slits in the end and put in little ribs. You can see the extent of the ribs. They are maybe a millimeter deep there. We put in those ribs and welded every intersection of rib to

sheet, and then threw the whole thing in nitric acid and to eat out the copper. You can see the ribs sticking out back there. This stack really looks nice.

(Transparency)

It is a very open structure. Light comes through it quite easily. We never made very careful measurements of the spacing uniformity, but it looked very good to the eye.

(Transparency)

The cold heat exchanger was made of copper fins soldered into a copper case, with little tubes going through in two planes. There are four tubes in each of these two bundles. That was all soldered together, and then we cut the tubes off flush here and here. There were some little water manifold spaces around the outside, so that each fin segment was cooled either by contact to two sets of water tubes or to one set of water tubes and the water-cooled case itself.

(Transparency)

Then that heat-exchanger assembly slipped into the conical reducer; the resonator is extending off into the distance there.

That is what the hardware looked like.

DR. GARRETT: Greg, was there a seal between the cold heat exchanger and the collar it fit in? What made the water seal when you shoved the heat-exchanger unit --

DR. SWIFT: That is a good question, by which I mean that we went through a few iterations on that, as we did on the hot heat exchanger.

(Transparency)

The original plan was that this would be all soft-soldered together and when we slipped this in, we were going to Cerro-solder this into the case, and then the water lines in and out would access these little annular spaces. That worked fine until we put the 30-bar helium pressure in.

The case sprung just enough under pressure to crack that Cerro-solder, and we had helium shooting into the water and making bubbly noises in the house water system.

(Laughter)

Then in desperation -- it was easy to get apart -- we actually milled a little O-ring groove right around that surface, around each of the two water header spaces. It is a hard O-ring to get into the groove and get it to stay in while you slide the thing in, but once it is in, it is good. If you are nervous about that, you add a second O-ring here as a backup. That is what we did.

(Transparency)

Here are some data. As I said, we had two stacks, the honeycomb stack, which was cheap and easy, and the spiral stack that we took great pains to make ourselves. The left half is a way of plotting the raw data. At this point, we only had one hot heat exchanger, so we were handicapped. We only stuck it in one end. One hot heat exchanger, one stack, and both cold heat exchangers.

You put it in, no load, turn up the heat, and observe what temperature the hot end runs at and what pressure amplitude you get in the resonator. That is the raw way of presenting the data.

Another way of presenting the data that we will talk more about in a little while is with things plotted in a dimensionless way, so that, instead of pressure amplitude, we plot pressure amplitude over mean pressure squared, and instead of heater power, we plot 10^4 times heater power over mean pressure, sound speed, and area, which is also dimensionless, and, instead of hot temperature, we plot hot temperature over cold temperature.

This is how well it does. You can look at the data on either the left side or the right side. In the limit of low amplitude or in the limit of low heater power, the measurements are in pretty good agreement with the design calculations.

At the high amplitude or high heater power extreme, there are some significant deviations -- 0.01 is a 10 percent pressure amplitude. That is where we tend to end up trying to operate in our Los Alamos projects. Some of the thermoacoustic refrigerators operate around 5 or 4 percent pressure amplitude. The square of 5 percent is about here. There is generally better agreement between calculation and measurement on some of the refrigerators that people have built. Those were some measurements we made at Los Alamos before we shipped the thing to Tektronix.

(Transparency)

Of course, they had their own purposes in mind. We really never got any great data under load, with the refrigerator attached. Charles Jin did at one point take one careful data point -- relatively careful; I think the uncertainties in everything here are about 10 percent -- well, not pressure amplitude. Whether he was really putting 490 watts to his load or 10 percent more or less, I do not know.

Heater power may be 5 percent uncertain. Hot temperature -- we had about a 100-degree uncertainty on the hot temperature when we were really moving heat, because the thermocouples were not right in the Nichrome ribbon; they were out on the case on the outside.

Anyway, the agreement is pretty good. There are 10 percent differences between experiment and calculation. Good enough for them to be pleased, anyway.

To continue on to the last part of the lecture, we have been looking at some data from that Tektronix thermoacoustic engine. At low pressure amplitudes, over to the left side of this graph, the measurements were in good agreement with the calculations.

(Transparency)

Out at high amplitudes, they were not so good. Give me some ideas. What did we leave out of the theory or what did we do wrong in the experiment? Where would you start if you were working for Tektronix and were faced with an extra 150° K, 300° F, hot temperature that ruins all your assumptions about what material you were going to be able to use on the hot end?

MR. HALLAJ: Unit losses are going up, and things like insulation are not going to be as affected.

DR. SWIFT: Thermal insulation around the hot end?

MR. HALLAJ: That is one loss, yes.

DR. SWIFT: I think I included that in that calculation of heater power.

MS. GRANT: Did you include viscosity?

DR. SWIFT: Viscosity is in there.

STUDENT: Did you include heat conduction through the stack parallel to the flow?

DR. SWIFT: Parallel to the flow, yes, in the metal of the stack, in the gas in the stack, and in the pressure vessel around that.

MR. HALLAJ: How about your assumption of laminar flow? I am not sure how high the velocities are, but eventually, I guess, you will start to get turbulence. The Reynolds number is very large.

DR. SWIFT: The Reynolds number in the stack is close to 1000.

MR. HALLAJ: So that is getting close.

DR. SWIFT: The Reynolds number out in the resonator is close to a million. That is not part of the acoustic approximation, so that is something that is definitely missing from the theory.

DR. HOFER: The Reynolds number requires a characteristic length. When we spit out these Reynolds numbers, we have to be --

DR. SWIFT: Yes. I was using, in both cases, the distance between plates or the diameter of the pipe.

DR. MIGLIORI: Have you looked at things with bigger exponents than two?

DR. SWIFT: In the energy equations? Can I write down higher-order terms?

DR. MIGLIORI: Yes, that is what I meant.

DR. SWIFT: In this acoustic approximation stuff, we throw away anything in the continuity and momentum equations that is second order, and we throw away anything in the energy equations that is beyond second order -- a lot of terms.

DR. HARGROVE: This may be saying the same thing, but nonlinear acoustic effects, and particularly enhanced portions, could be a diversion of energy to higher harmonics.

DR. SWIFT: Okay, harmonics.

DR. GARRETT: Did you have a mike in there where you could look at the waveform shape, Greg?

DR. SWIFT: Yes.

DR. GARRETT: How did it look?

DR. SWIFT: It looked about like this, flat-bottoms, "pointy" tops. Harmonics. It is not in the theory.

DR. GARRETT: Not sinusoidal?

DR. SWIFT: Yes, that is right.

DR. MARSTON: In your early equations, you had this curious term that involved the gradient of the shear viscosity. Did you keep that in your code or not?

DR. SWIFT: No. At the first order, that does not show up. Not keeping it is consistent with all the other approximations that we make. But that is something that I would add to this topic.

DR. GARRETT: Something we do not have in refrigerators where the temperature span is more modest -- is it possible that there is a longitudinal variation in the transport properties as you approach plates, that the viscosity is not a constant?

DR. SWIFT: I think that is what Dr. Marston just asked. In the governing equations, there are derivatives that look like one derivative of μ grad u , or something like that. To first order, you have a u_1 and you pull μ out of the derivative, but to higher order you cannot do that.

DR. CHAMBERS: How good is the coupling between the stacks and the heat exchangers? Are there any breaks in there, or gaps?

DR. SWIFT: That is a good question. I am glad you brought that up. Somebody else asked that when we took our break. Let me get a diagram up here so we can talk about that.

(Transparency)

These do not touch, exactly. There is probably a tenth of a millimeter gap here, and the gap here is as big as it was. The gas-displacement amplitudes are a centimeter or two. So it is not metallic thermal contact that is carrying the heat from here to here. It is the gas moving back and forth. It really does not matter whether they touch or not.

MR. HALLAJ: How fast do the particles travel in that, what is the velocity?

DR. SWIFT: Do you remember my viewgraphs that had the colored hand traces of velocity and pressure? Eighty meters a second in the middle of the resonator? It is more like 5 or 10 meters per second down here.

MS. GRANT: Is the heat exchanger keeping up?

DR. SWIFT: Heat exchanger bottleneck? So it could be that we cannot pump water through fast enough or the heat cannot get out from the middle of the fin to where the water is, or something like that. That is not in the calculations.

DR. MIGLIORI: Two other things. Have you looked at radiation losses, thermal radiation?

DR. SWIFT: Blackbody radiation?

DR. MIGLIORI: And if the plates vibrate, the flow impedance goes like the spacing squared -- weird mechanical problems. The engine does not look the way you think it looks when it is operating.

MR. HALLEJ: That could be even more dramatic if you are occluding the space between -- if you get the boundary layers. You said you had 80 meters per second, which is a good fraction of the Mach number, and for some temperatures that spacing that you have for flow may not be what you think it is. You have a couple of boundary layers in there. It can occlude the space.

DR. SWIFT: Yes.

DR. MIGLIORI: So what is the answer?

DR. SWIFT: This is a pretty good list. It would take a lot of people a lot of time to work through all this. What are we going to do? Honestly, when you think about it, we really do not know a whole lot about high-amplitude operation in these thermoacoustic devices.

Let me see if I can add any ideas here: roughness on the surface of the plates. You might imagine that if the roughness height is of the order of even a tenth of a viscous penetration depth, so much for our pretty pictures of the oscillatory flow. There is another thing to worry about.

DR. GARRETT: Entry effects. It is entering and leaving 800 times a second. It has to go from this turbulent mess in the mega-Reynolds number area to this nice, ordered flow. We know from pipe flows that there is a distance it takes to establish laminar flow, even when there should be laminar flow.

DR. SWIFT: Yes. So in the first several millimeters here, our theory is no good even to first order.

FACULTY: Oscillatory flow does not make it harder to deal with?

DR. GARRETT: It probably does not make it any harder to deal with. It just means you cannot go to the literature, right?

DR. ATCHLEY: It makes it harder to deal with when the heat exchanger is not the same length --

(Simultaneous discussion)

DR. RASPET: There is something intrinsic in here, and that is that, for various arguments, you assume pressure is a plane. These high amplitudes -- would the pressure not be the same value all the way across? You are probably getting large enough flows to have significant axioradial flows.

DR. KEOLIAN: These guys are sort of mumbling that it may be part of the heat-exchange bottleneck, but the thermal contact between the gas and the heat exchanger may have a weird --

(Simultaneous discussion)

DR. HOFER: If you think about nonuniformity of the stack channels and the fact that they may be thin and flimsy, and there is, as has already been mentioned, vibration, but you can also throw in acoustic streaming. This would be a higher order term, but acoustic streaming is dependent on nonuniformity or the fact that it is vibrating.

Or if you think it is vibrating, it might be pure friction. You have these little welds in your honeycomb, and you have friction there. Or maybe the whole stack is hammering back and forth in there, with a lot of friction being generated.

DR. SWIFT: Hammering. We might include hammering under the general category of experimental "blunders."

(Laughter)

It is really complicated.

DR. MIGLIORI: If you are not going to give us the answer, we are not going to help you any more.

DR. SWIFT: I was going to ask a second question now: What can we do experimentally to figure out which one of these is the culprit? That could take us the rest of the day.

(Transparency)

I want to show you just one thing that we have gotten in the habit of doing that has been a little bit of help in trying to think about these complicated issues. That is thinking in terms of dimensionless groups and similitude and stuff like that.

You know this classic problem of a sphere moving in a viscous fluid, like a baseball or a physicist's version of a person dropping through air. You would think that the force on that moving object would be some function of velocity, density, viscosity, and radius of the sphere. If you did not think about it any more carefully than that, you could fill up a handbook with tables of what was going on, to fully characterize the force on the moving sphere.

But once you appreciate the power of making dimensionless groups out of things, such as the Reynolds number, you find out that you can write this as a function of only one variable, the Reynolds number. Now the table, instead of being a big, fat handbook, turns into one page.

Furthermore, if you make some measurements that do not follow this simple law, then you know you have other physics in the problem that you have to start thinking about. What is it? Maybe then you think about compressibility and Mach number and stuff like that.

(Transparency)

If you think through this similitude stuff for thermoacoustics, with some generic thermoacoustic apparatus drawn up here, with a bunch of powers like acoustic power and heat, and a bunch of dimensions like length and x_j , where j probably goes up to 100 or something, and all the gas properties that matter -- the ratio of specific heats and the sound speed and the thermal conductivity and the viscosity -- and the exponents that give the temperature dependence of viscosity and thermal conductivity, and you think about how to make dimensionless groups -- there is not a unique set, but one particularly nice set is what is shown down here at the bottom.

If you take the frequency at which this engine springs into spontaneous oscillation and scale that by length and sound speed, you get a dimensionless number. If you take the temperature

at each point, scale it by some reference temperature -- we always use the room-temperature heat exchanger as a reference, but you do not have to do that -- if you normalize all the pressures by the mean pressure, all the velocities by the sound speed, all the powers by mean pressure, sound speed, and cross-sectional area of the stack, and make some dimensionless groups over here, including the Prandtl number and including the thermal penetration depth divided by the pore size in the stack, then maybe you can make some progress in trying to figure out which of these guys is to blame and which is not, and how to deal with it.

(Transparency)

For instance, we had an engine at one point that looked like this. We put in heat to a hot heat exchanger with electric resistance heaters. If we put in enough heat, the thing would spring to life and oscillate in steady state. We could measure the hot temperature and measure the pressure amplitude and so on.

If you looked at the harmonic content of those pressure oscillations, here is the 2f component normalized by mean pressure, and here is the square of the fundamental normalized by mean pressure. It is a straight line. Okay, that sounds like good old nonlinear acoustics. No big deal. The second-order term ought to be the square of the first-order term, and it is.

But all these different gases we tried -- helium, neon, argon, nitrogen, in the same apparatus -- all the different gases fell on the same straight line. We changed mean pressure, and changed mean pressure again for argon. I think the mean pressure here might have been three times the mean pressure here. It is still on the same straight line.

(Transparency)

Does that, plus looking at these dimensionless groups, tell us anything about some of these? We are on harmonics now. We are talking about harmonics. What do we now know that the harmonics do not depend on? When we switched from pure gases to a gas mixture, we changed the Prandtl number and not anything else in this list. It does not look as if the harmonic content depends on Prandtl number -- progress.

When we change from monatomic gases to nitrogen, a diatomic gas, harmonic generation does not seem to depend on gamma. When we change mean pressure, going from this column to this column to this column, we change thermal penetration depth, because it has the density in it, and the density depends on pressure. Harmonic generation is not depending on that. Maybe we can make some progress.

(Transparency)

Some more data from the same engine: again, fundamental pressure amplitude over mean pressure squared; 10 percent amplitude, zero; hot temperature over cold temperature; heater power normalized by mean pressure, sound speed, and area.

We tried three different gases. In each case, we selected the mean pressure so that the thermal penetration depth was 89 percent of the pore size in the stack. Similitude says that if we plot with these dimensionless groups, it should not matter which gas we have. Indeed, it looks like it does not matter which gas we have. That is reassuring.

This deviation from our acoustic approximation calculations is included in the similitude analysis. I think that means that it does not have anything to do with plate vibration, because we do not have Young's modulus of the plates or anything in the analysis.

DR. MIGLIORI: You worry about affecting just the viscous term only. If they get closer together you get d squared for the flow impedance, and it is a smaller effect. That should have been in there.

MR. HALLEJ: You could group these by what nondimensional number they depend on and just kind of treat them as --

DR. SWIFT: Yes. You can use this kind of analysis to eliminate whole blocks of problems. You can switch to a different set of mean pressures for your three gases so that you change the ratio of penetration depth to pore size significantly. We made a big change here.

The deviation from the acoustic approximation has something to do with this variable δ over h . All right, we are making progress now.

(Transparency)

Here is one for the people in the back row. In this same engine that we have just been talking about, we accidentally stumbled upon a regime of operation that was unstable -- usually with these engines, you put some heat to them and they steadily oscillate at some amplitude they choose. But this one, we had argon in it at one time, and we had a particular mean pressure and a particular heater power, and the darn thing would not operate steady state.

It would shut itself off for 20 minutes, and then it would run for 5 minutes, then it would get quieter and quieter and shut off again, completely off, for 20 minutes, and come back on again. We thought, well, what is going on here?

All right, we know what to do. We will switch gases and see whether the origin of that is blackbody radiation or -- what is going on here? So we switched gases, and the map of the instability region on this dimensionless plot turned out to be the same. That makes it sound like this effect has a lot to do with gas dynamics, nonlinear effects in the gas equations or something like that. Something included in our similitude analysis.

Here is the question for you guys in the back row. The period with which the thing would shut off and turn back on again, did not scale properly according to the sound speed of neon relative to the sound speed of argon.

DR. MIGLIORI: You lost me completely at that point. You have thought about this enough, so you have figured out what effect and what dimensionless number you can connect up.

How did you connect up the sound velocity of the gas, the sound speed of the gas, with this instability in the 20 minutes?

DR. SWIFT: Frequencies always scale with a length and a sound speed. Pick one length, whatever length you want, and do that throughout the whole analysis of a given problem. Sound speed is sound speed.

DR. MIGLIORI: You must have tried it, then, with viscous and thermal penetration depths or something, as well.

DR. SWIFT: We kept the ratio of penetration depths to those dimensions the same, so it does not matter if you switch to a different length, like the penetration depth.

DR. MIGLIORI: What is r ?

DR. SWIFT: Pore size in the stack, hydraulic radius.

DR. MIGLIORI: So that keeps all of the transverse heat flow time constants the same.

DR. SWIFT: Yes. It is really pretty powerful.

DR. BASS: How do the thermal relaxation times compare to the thermal relaxation times if the stack is a solid or a tube? Twenty minutes is a pretty long time.

DR. SWIFT: I think you are on to something there. The only place you are going to get 20 minute time constants is in the heat capacity of some big hunk of metal in there. That has to be what is going on.

DR. KEOLIAN: I do not see how you get to that from this similitude analysis.

DR. SWIFT: The similitude analysis is based on the gas equations, the momentum equation, the continuity equation, energy equation, all that stuff. It has the gas variables in it. It does not have the Young's modulus of steel in it. If you have an effect you do not understand and you make a change that ought to work with your similitude analysis and it works, then it is circumstantial evidence that the physics behind your effect is in the variables that you started with, in the equations that you started with.

DR. KEOLIAN: In the gas.

DR. SWIFT: In the gas. But if you make a change and it does not work at all, it is like the old problem of dropping spheres through a viscous fluid. If it does not scale properly with similitude, you may have gotten yourself into a new regime where you need the Mach number in the problem of dropping a sphere through a viscous fluid, and you had only the Reynolds number before. Or you need the surface roughness of the ball. You have a golf ball instead of a smooth sphere.

DR. KEOLIAN: But if you try it for that whole list of 15-something dimensionless variables and it does not scale with those guys, then you --

DR. SWIFT: It is not an ironclad approach to anything, but it can be helpful.

MR. HALLAJ: None of those non-dimensional variables contained the heat capacity of the metal. A test might be to build one out of aluminum and then build another out of --

DR. SWIFT: Yes.

DR. MIGLIORI: Did you conclude anything from all of this?

DR. SWIFT: No. They are just questions.

DR. MIGLIORI: I mean this particular point. You established that none of the normalized time scales explained this. Is that right?

DR. SWIFT: Yes.

DR. MIGLIORI: So your assumption now is that you can look for other physics.

DR. SWIFT: The responsible party has to be some physics that we had not thought of when we set up that table of dimensionless groups. Dr. Bass pointed out one such candidate that seems plausible.

DR. MIGLIORI: When the engine shuts off, all the temperature gradients instantly change very rapidly, don't they?

DR. SWIFT: Yes.

DR. MARSTON: Is there an obvious reason to see why the similitude seems to capture the stability boundary, but not the frequency? In other words, why do you get one and not the other?

DR. SWIFT: I do not know. I am making a big deal out of this similitude stuff because, no matter what you are working on in your research, you ought to be thinking about it. There may be something for each of you. For example, it may be that in sonoluminescence there is some way to make dimensionless groups that will tie a couple of things together or exclude something else.

DR. KEOLIAN: I have another question that is related. If we all work with the same dimensionless groups, is it valuable for comparing radically different engines? If I do something and you do something, and I say mine is and you say yours is not? Is that meaningful?

DR. SWIFT: Yes, I think it really is. I think you want to know, for instance, whether you have the same ratio of thermal penetration depth to stack pore size. Those are things we already do as a matter of course, to a large extent.

Dr. Hofler and I were talking about this earlier in the week. He has a particular engine that has room temperature on the hot end, liquid nitrogen on the cold end, with neon at half an atmosphere. There is a whole family of engines that are strictly similar to that, that operate between 1200° K and 300° K, the same temperature ratio, and that have scaled dimensions and different gases.

One of those engines might be something that you are interested in, and you have not thought about it because you have not thought that half-bar neon is so similar to 5-bar helium in a different temperature or size range.

(Transparency)

One more use of similitude is in the building of scale models. In the airplane industry they do not usually build the 747 without building a little scale model and testing it in a wind tunnel. It is cheaper to make scale models and test them in wind tunnels.

At Los Alamos, we are involved in a project to make a big version of that Tektronix device, a heat-driven cryogenic refrigerator. In this case, the heat source will be burning natural gas, and the object of the game is to liquefy natural gas. Ultimately, we want to end up with a thermoacoustic resonator about that long, with stacks in both ends and so on.

But our lab is not this big. Similitude tells us how to make a scale model that is completely meaningful, that captures all the things we do not understand within the gas equations.

(Transparency)

It turns out, if you want to make a half-scale model of something that at full size had 30 bar helium, so on and so on, you build at half-size, you put in argon at a slightly reduced pressure, and everything turns out to be the same. The ratio of penetration depth to pore size is the same. The temperatures are the same. The pressure ratios are the same. The powers normalized by p mean sound speed area are the same. That means that the actual powers, in watts, go down by 16 in the scale model, which is really nice because you do not need to handle 100-kilowatt powers in your little laboratory.

(Transparency)

David Gardner has been involved in this. The first thing that our industrial collaborator is going to build is not that big thing that I showed you the cartoon of on a truck, but it is something that is twice as big as this apparatus that is sitting behind Dr. Gardner here, with thermoacoustic engines, stacks, and heat exchangers under this insulation space, and then the resonator extending down to here.

The first thing that these industrial guys always ask is, can you make it cheaper or can you make it smaller? They wanted to make it smaller by making it shorter, by just making some cuts in the resonator and some miter bends and bringing the resonator back around. We can answer that kind of question in short order. You can probably guess that if you had Reynolds numbers of a gazillion and Mach numbers of 0.1 down there, a corner like that is a bad idea.

(Transparency)

It is not clear whether that might not be a better idea. Anyway, scale models are a very important use of similitude.

I have three summary viewgraphs and then we will call it quits.

(Transparency)

This is what we know how to do best calculationally. This is what was pioneered by Nicholas Rott 20 and 30 years ago. You start with what we always start with in acoustics, an

equation of state for the material -- in this case, gas -- a version of Newton's law that is appropriate to us, a version of mass conservation that is appropriate to us, and some heat-transfer equation.

We apply the acoustic approximation, let somebody else do the integrals perpendicular to the x direction once and for all, and then we can do a numerical integration along x and find out just about everything we want to know within the acoustic approximation. There is no longer a whole lot of debate about the details of this process in the acoustic approximation.

But there is going to be a lot of debate and a lot of work to be done on things such as we came up with here, and I hope a lot of effort at inventing really fun gadgets and really important applications for these devices.

(Transparency)

Cost and efficiency are always the things that control what applications get developed. Why can't thermoacoustic systems have Carnot's efficiency? I wish they did, but they do not. Here is a list of things that contribute to inefficiency. Very roughly, they are all equally important. There is nothing on this list that you can neglect *a priori* when you walk up to a thermoacoustics problem.

Topping the list are the thermal relaxation losses within a thermal penetration depth that are what make the whole effect work. The whole reason that those little ovals in the indicator diagram and the animation on the pV plot are ovals and not reciprocating lines is because of heat transport across a thermal penetration depth. That is irreversible heat transfer, and we are stuck with it.

Viscous dissipation is important. It is hard to get away from that in ideal gases because the viscous penetration depth is just about as big as the thermal penetration depth. All the surface area in here is experiencing viscous shear.

Conduction of heat along x -- it is just the plain, old dumb conduction from here to there -- is a problem that you have to watch out for.

These first three effects take place not only in the stack, but also elsewhere. Out here in the resonator, there are viscous losses. This surface of the resonator has thermal relaxation losses. There is conduction of heat along x in the pressure vessel that you build around that part. There are bottlenecks in the heat exchangers that eat into your efficiency a little bit.

Transducers -- we have not even talked about transducers. I have been focusing on a thermoacoustic engine that does not need a power transducer either to put in power or to take it out. But if you are in the refrigerator business, you have to worry about the efficiency of converting electric power into acoustic power.

Then there are all of these things that we do not really understand yet. About one-third of Carnot is typical of current designs. About one-fifth of Carnot is the best that anybody has done so far. It always takes a while to build something, and we have gotten smarter in the design since the last time we built something.

STUDENT: The first source you have to have. If you could do away with all the others, what is your feel for how close you could get?

DR. SWIFT: That is a good question. If you could get rid of all the other inefficiency mechanisms and only leave number one, how well could you do? I think you could do arbitrarily well in efficiency if you were willing to sacrifice power per unit volume.

You remember in the final animation there was one frame where I had the temperature gradient in the stack just matching the temperature profile that the gas experienced as it went back and forth. In the zero-viscosity picture, there are not any losses there. There is no power either. The area of the pV oval is zero.

I think, as you approach that line arbitrarily closely from either side -- if you are talking about either a refrigerator or an engine -- your irreversibilities would then go to zero, so you would have Carnot efficiency. But at the same time, the power is going to zero, so you would have to make the thing bigger and bigger to accomplish a given task.

MR. POESE: (Inaudible.)

DR. SWIFT: But he wanted to get rid of that and just look at number one.

MR. POESE: But the bigger you make it, the harder it is to get rid of all those inefficiencies.

DR. SWIFT: Yes.

(Transparency)

A final summary viewgraph. Keep an eye on similitude, no matter what you are doing. We found it to be a really powerful, thought-provoking technique, and maybe you will, too.

Thanks a lot.

DR. LEVY: How about these Third World applications for these refrigerators? How is that progressing? I saw some of your pictures a few years ago.

DR. SWIFT: Dr. Levy was asking about applications of thermoacoustics in developing countries or undeveloped countries.

A couple of years ago, we were contacted by a guy in Bangladesh who recognized the need to keep vaccines cold in rural medical centers, where there is no electricity. He said the way they do it now is to have guys on bicycles with ice chests scurry back and forth from the cities to bring ice to the rural medical centers.

We cooked up a design for a combustion-powered refrigerator that would use 60 psi air as the working substance and burn kerosene to keep the medicines cold. He was moving ahead on that and then he got a grant from one of the international development agencies to complete his Stirling engine rice mill project, which was very dear to his heart because he had worked on it for 10 years before and it fizzled due to lack of funds.

He and his people are back on that full-time, with the thermoacoustics project on the back burner. They have all these nice Stirling rice mills now out in the countryside that are efficient enough to mill the rice using the combustion of the hulls of the rice as a power source.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE		3. REPORT TYPE AND DATES COVERED Final 01 Feb 96 - 31 Jan 97	
4. TITLE AND SUBTITLE Proceedings of the 1997 Physical Acoustics Summer School: Volume I: Transcripts				5. FUNDING NUMBERS PE 61153N G N00014-96-1-0033	
6. AUTHOR(S) Henry E. Bass					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Mississippi NCPA University, MS 38677				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research ONR 331 800 North Quincy Street Arlington, VA 22217-5660				10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; Distribution unlimited				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Volume I of the Proceedings of the 1997 Physical Acoustics Summer School contains verbatim transcriptions of the lectures given at the school. Volume II of the proceedings contains copies of the transparencies used during the lectures and Volume III contains copies of the background materials sent to student participants prior to the summer school.					
14. SUBJECT TERMS Physical Acoustics, Molecular Acoustics, Acoustics Demonstrations, Resonant Ultrasound Spectroscopy, Sonoluminescence, Nonlinear and Atmospheric Acoustics, Sensor Physics, Thermoacoustics				15. NUMBER OF PAGES 368	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT		